EXHIBIT 2
[FILED UNDER SEAL]

UNITED STATES DISTRICT COURT

EASTERN DISTRICT OF TEXAS

SHERMAN DIVISION

The State of Texas, et. al. Plaintiff, v. Google LLC, Defendant.	Case No: 4:20-cv-00957-SDJ

Expert Report of Matthew Weinberg

6/7/2024

Matthew Weinberg

Table of Contents

I. INTRODUCTION

A. Assignment

1. I have been engaged by the counsel for the Office of Texas Attorney General, on behalf of all Plaintiff States in the case, to analyze several changes made to the auctions for online display ads[1] under the conducts undertaken by Google. In particular, I have been asked to opine on:

a. How the conducts undertaken by Google (including Dynamic Allocation, Dynamic Revenue Sharing, Project Bernanke, Unified Pricing Rules and Reserve Price Optimization) change the auction procedure from an auction theory perspective, and

b. How the resulting changes to auction procedures impact outcomes (such as revenue, fill rate, price, payoff, win rate) for various actors in the auctions for online display ads (including publishers, advertisers, exchanges, ad buying tools).

2. I am being compensated for my services in this matter at my hourly rate of $500. I have been supported by a research team who worked under my direction. The opinions and conclusions in this report are my own. My compensation is not dependent on my opinions, my testimony or the outcome of this case.

B. Qualifications

3. I am an associate professor of computer science at Princeton University, where I have taught since 2017. I am also the associate director of Princeton University's Center on the Decentralization of Power (DeCenter) and affiliated with Princeton's Center for Information and Technology Policy (CITP), Program in Applied and Computational Mathematics (PACM), and Bendheim Center for Finance. I received my B.A. in Mathematics from Cornell University in 2010, and my Ph.D. in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology in 2014.

4. My field of expertise is Algorithmic Mechanism Design, which is the study of algorithms (such as ad auctions) that involve economic incentives (such as those of publishers, exchanges, ad buying tools, and advertisers). I have published over 50 papers across venues within both

[1] Online display ads a specific type of targeted online advertisements that appear on a publisher's website along with their content.

computer science and economics. My works have received the top dissertation award and the best full paper award from ACM SIGecom, the interdisciplinary research community studying Economics and Computation, and a test of time award from IEEE Foundations of Computer Science, the research community studying Theoretical Computer Science. For my research and teaching, I have also received a Sloan Foundation Fellowship, an NSF CAREER Award, Princeton's President's Award for Distinguished Teaching, a Phi Beta Kappa Teaching Award, and an Engineering Council Lifetime Achievement Award.

5. As I state below, the market for online display ads mainly relies on auctions to conduct the sales of display ads. As a result, as my qualifications above show, I am qualified to opine and conduct analysis on the market for online display ads.

6. My CV is attached in Appendix A and contains a list of my publications in the past ten years.

C. Case Background

7. A coalition of 16 states (Texas, Alaska, Arkansas, Florida, Idaho, Indiana, Kentucky, Louisiana, Mississippi, Missouri, Montana, Nevada, North Dakota, South Carolina, South Dakota, and Utah) and the Territory of Puerto Rico (Plaintiff states), led by the State of Texas, filed a lawsuit against Google LLC (Google), claiming Google violated federal and state antitrust laws and other state laws, in Google's conduct in the online display advertising market. I was retained in September 2021 to provide expert analysis and opinions on behalf of all the Plaintiff states.

8. I also understand that the Plaintiff States have the opportunity to submit rebuttal expert testimony and reports and that I may be asked to evaluate the reports and opinions offered by Google's experts in connection with those rebuttal reports.

9. A list of all documents referred to in this report and relied upon by me in forming my opinions in this case is attached as Appendix B. I have reviewed, signed, and complied with the Confidentiality Order entered in this case. My supporting team has also read, signed, and complied with the Confidentiality Order entered in this case. I have also reviewed the Stipulation and Order regarding Expert Discovery in this case.

10. I understand that document productions are ongoing in this case and that additional relevant documents may be produced in this case by Google and third parties right before and after I issue this report. I may, and reserve the right to, review and rely on additional documents

in conducting my work and forming my opinions in this case. I reserve the right to supplement or amend this report if my opinions change or require supplementation as a result of my ongoing review of documents.

D. Summary of Opinions

11. I have analyzed each of the conducts undertaken by Google and assessed how they change the auction procedure and how these changes affect the auction outcomes.

12. I conclude that:[2]

a. Google's implementation of Dynamic Allocation led to higher win rate and higher revenue for AdX[3] as well as lower win rate and lower revenue for non-Google exchanges. Furthermore, Enhanced Dynamic Allocation led to an increase in win rate and increase in revenue for AdX and reduced the value of direct deals for advertisers. Reducing the value of direct deals for advertisers would decrease the revenue earned by publishers via direct deals.

b. Header bidding improves publisher outcomes relative to the waterfall approach (with or without Dynamic Allocation and Enhanced Dynamic Allocation) and it can generate higher revenue for publishers compared to Exchange Bidding.

c. Unified Pricing Rules likely lead to lower revenue for the publishers. It also can lead to better win rate and revenue for Google's ad exchange AdX as well as Google's ad buying tools and lower the win rate and revenue for rival exchanges and ad buying tools.

d. Under the Dynamic Revenue Sharing (DRS) conduct,

    i. Dynamic Revenue Sharing version 1 (DRSv1) increased AdX win rate and revenue and decreased non-AdX exchanges' win rates and revenues, compared to no DRS,

[2] These conclusions discuss the isolated impacts of the conducts, and throughout the report I also provide conclusion regarding the interactions between the conducts.

[3] AdX is Google's ad exchange.

ii. Dynamic Revenue Sharing version 2 (DRSv2), in comparison to both no DRS and DRSv1, decreased advertiser payoff,[4] increased AdX win rate and revenue, decreased non-AdX exchange's win rates and revenues, and may also have decreased publisher revenue.

iii. Truthful Dynamic Revenue Sharing increased AdX win rate and revenue and decreased non-AdX exchange's win rates compared to no DRS,

iv. Google concealed information that is vital to advertisers and important to publishers by concealing DRSv1 from them.

e. Projects Bernanke and Global Bernanke did not affect GDN[5] advertisers and could increase some publishers' revenues while decreasing others. However, these projects also led to a lower win rate for non-GDN ad buying tools and advertisers that used those ad buying tools. Furthermore, Project Bernanke and Global Bernanke led to an increased win rate for GDN buyers (without improving GDN advertisers' payoffs), which leads to an increased win rate and revenue for GDN.

f. Reserve Price Optimization leads to higher revenue for Google's ad exchange AdX, and lower payoff to advertisers. It could also lead to lower payoff for some publishers. The negative effects of Reserve Price Optimization to advertiser payoff, and possibly some publishers' revenues, is due to (a) Google's concealment of the conduct during its initial rollout, and (b) barriers to publishers effectively setting reserve prices to optimize their revenue even after Google announced the conduct.

E. Methodology

13. Throughout the report, I am going to apply the mathematical principles, results and insights that stem from the canonical auction theory and game theory literatures. These methods of auction analysis are commonly accepted by researchers and practitioners across many different fields such as economics, computer science, and mathematics. Furthermore, these tools are commonly accepted by researchers and practitioners for the analysis of the market at hand, online display ads.

[4] I use the term "advertiser payoff" to refer the difference between the advertiser's value for the impression and the amount they pay for the impression.
[5] GDN refers to Google Display Network, Google's ad buying tool for small advertisers. Its current name is Google Ads.

14. I overview relevant definitions, results and principles regarding the auction theory fundamental necessary to analyze the case at hand. I will use those stated results and principles of auction theory to analyze the facts of this case.

II. RELEVANT CONCEPTS IN AUCTION THEORY

15. In this section, I introduce concepts in auction theory relevant in understanding online display advertising and the associated tools. I provide an overview of what an auction is, important components of auctions, common formats used to conduct auctions, strategic behavior in auctions, and the role information plays in auctions.

A. Auctions

16. Most people have seen dramatic illustrations of auctions on television, where a lively auctioneer for an artwork solicits bids from the audience and participants repeatedly raise their paddles to bid higher and higher. This specific auction closes when paddles stop flying and the item is sold to the highest bidder.[6] Auctions can vary in complexity, specificity/concreteness, and format. Common to all auctions is a process that transfers items from seller to buyers, often with payments from buyers to sellers.[7] Auction theory studies these processes and their outcomes, such as the allocation of the items and associated payments. I explain the relevant auction formats to the case at hand below.

1) The First- and Second-Price Auctions

17. I focus on **single-item auctions**, where the seller has a single item for sale and there are multiple bidders interested in the item.[8] A single-item auction can be held with different rules. This is referred to as the auction format. One such single-item auction format is a **sealed bid auction**, where the auctioneer solicits a single bid from each bidder and directly decides from these bids to whom to award the item and how much to charge. The term "sealed bid" refers to the fact that each bidder submits a single bid to the auctioneer, in a manner so that other bidders cannot see, and has no further communication with the auctioneer.[9] Every sealed bid auction has two

[6] This particular format is called an *English auction*, also known as ascending price auction.
[7] Vijay Krishna. "Auction Theory" (2009). Academic Press. (Chapter 1.)
[8] There are many auction types other than single-item auctions, but auctions for online display advertisements considered in this report are single-item auctions, so I focus on this specific type.
[9] Sealed bid auctions are also called direct revelation auctions. The term "direct revelation" refers to the fact that each bidder "directly reveals" a bid to the auctioneer, rather than participating in a possibly complex process (i.e., there is no negotiation, no sequential offers, *etc.*).

components: An **allocation rule**, which determines who (if anyone) gets the item, and a **payment rule** which determines how much (if anything) each bidder pays.[10]

18. The most common auction formats used today are called first-price and second-price auctions. In a **first-price auction**, the auctioneer solicits a bid from each bidder, the highest bidder wins, and pays their bid to the auctioneer. The name "first-price" refers to the fact that the payment is equal to the highest bid.[11, 12] Using the terminology above, first-price auctions are a type of sealed bid auctions where the allocation rule awards the item to the highest bidder and the payment rule charges the winner their bid and does not charge other bidders. In a **second-price auction**, the auctioneer solicits a bid from each bidder, the highest bidder wins and pays the highest other bid to the auctioneer. The name "second-price" refers to the fact that the payment is equal to the second-highest among all bids.[13] In other words, second-price auctions are a type of sealed bid auctions where the allocation rule awards the item to the highest bidder and the payment rule charges the winner the highest other bid and does not charge other bidders.

19. The first-price auction and second-price auction have the same allocation rule (they both award the item to the highest bidder), but different payment rules (one charges the winner their bid, the other charges the winner the highest other bid). To concretely set the ideas, take the following example.[14] Imagine there are five bidders, who submit bids of $1, $8, $3, $5, $2, respectively. Then in both the first- and second-price auction, Bidder Two (the highest bidder) wins. In the first-price auction, Bidder Two pays $8 (their bid). In the second-price auction, Bidder Two pays $5 (the second-highest bid).[15, 16] Figure 1 below illustrates this example.

---

[10] The resulting payment of an auction is called the clearing price.

[11] First-price auctions are also called, pay-your-bid auctions. The term "pay-your-bid" refers to the fact that the winner pays their own bid.

[12] To be more mathematically precise, imagine there are n bidders labeled 1, …, n. Then each bidder i submits a bid b_i to the auctioneer. If i^* denotes the bidder $i^* := \mathrm{argmax}\{b_1, \ldots, b_n\}$ who submits the maximum bid, then i^* wins the item and pays b_{i^*} . The notation ":=" denotes that we are defining the left-hand side (i^*) to be equal to the right-hand side ($\mathrm{argmax}\{b_1,\ldots, b_n\}$). The notation $\max\{b_1,\ldots, b_n\}$ refers to the maximum number among $\{b_1, \ldots, b_n\}$. The notation $\mathrm{argmax}\{b_1,\ldots, b_n\}$ refers to the subscript i for which $b_i = \max\{b_1, \ldots, b_n\}$.

[13] To be more mathematically precise, imagine there are n bidders labeled 1, …, n. Then each bidder i submits a bid b_i to the auctioneer. If i^* denotes the bidder $i^* := \mathrm{argmax}\{b_1,\ldots, b_n\}$ who submits the highest bid, and i_2 denotes the bidder $i_2 := \mathrm{argmax}\{b_1,\ldots, b_{i^*-1}, b_{i^*+1},\ldots, b_n\}$ who submits the second-highest bid, then i^* wins the item and pays b_{i2}.

[14] I will employ this example throughout this section to illustrate ideas. This example will be referred to as the "leading example."

[15] In the more mathematical notation, $i^* = 2$ (the highest bidder is Bidder Two), b_{i^*} = $8 (Bidder Two's bid is 8), $i_2 = 4$ (the second highest bidder is Bidder Four), and $b_{i2} = 5$ (Bidder Four's bid is $5).

[16] A relevant question here is what happens when there are ties. Imagine there are five bidders, who submit bids of $1, $8, $8, $5, $2, respectively. Then there are multiple "highest" bids, and we must somehow break the tie. It is common to break ties randomly (i.e., pick randomly between Bidder Two and Bidder Three) or to break according to some arbitrary tie-breaking rule (i.e., pick Bidder Two because they come first in the ordering). In both the first- and second-price auction, no matter how ties are broken, the winner is either Bidder Two or Three (a highest bidder). In the first-price auction, the winner pays $8 (their bid). In the second-price auction, the winner also pays $8 (the

Figure 1: Clearing prices differ between first- and second-price auctions, but the winners are the same bidders

	Winner	Payment
First-Price Auction	Bidder Two	$8
Second-Price Auction	Bidder Two	$5

20. Although the first- and second-price auctions share the similarities of directly soliciting bids and awarding the item to the highest bidder, they also differ in important ways, since the first-price auction charges the winner their bid, whereas the second-price auction charges the winner the highest other bid.

second-highest bid). Importantly, note that the "second-highest bid" always refers to "the highest remaining bid after removing the winner" ($8). It is not the "second-highest number among the submitted bids (which would be $5 in this case)." In the more mathematical notation, i^* is either 2 or 3, b_{i^*} = $8, i_2 is either 2 or 3 (whichever one i^* is not), and b_{i2} = $8.

21. In a single-item auction,[17] the **minimum bid to win**[18] of a bidder is equal to the smallest bid they could submit and still win the auction.[19] This concept can be illustrated with the leading example. Imagine there are five bidders, who submit bids of $1, $8, $3, $5, $2, respectively. Because the first- and second-price auctions have the same allocation rule, minimum bids to win are the same in both auctions for each bidder. In this example, Bidder Two's minimum bid to win is $5 (if Bidder Two submits any bid strictly larger than $5, they will win and if Bidder Two submits any bid strictly less than $5, they will lose). All other bidders have a minimum bid to win of $8 (they would need to submit a bid of at least $8 in order to win). Importantly, the minimum bid to win of each bidder is the highest other bid. Furthermore, observe that Bidder Two's bid ($8) exceeds their minimum bid to win ($5), and therefore they win. All other bidders' bids ($1, $3, $5, $2) fall below their minimum bid to win ($8), and therefore they lose.[20]

[17] More specifically, this applies to only monotone allocation rules. The allocation rule of a single-item auction is monotone if increasing the submitted bid, while keeping bids of all other bidders fixed, cannot cause the bidder to switch from winning to losing. To be more mathematically precise, an allocation rule is monotone if for all bidders i, all possible bids $b_1,\ldots, b_{i-1}, b_{i+1},\ldots,b_n$ of the other bidders, and all possible bids $b<b'$ of Bidder i: if Bidder i wins on bids $b_1,\ldots, b_{i-1}, b, b_{i+1},\ldots,b_n$, then Bidder i also wins on bids $b_1,\ldots, b_{i-1}, b', b_{i+1},\ldots, b_n$. Observe that "highest bid wins" is a monotone allocation rule. If a bidder is currently winning, then they must have the highest bid. That is, $b_i \geq b_j$ for all j that is not equal to i. If that bidder increases their bid, then that bidder must still have the highest bid. That is, the new bid $b'_i > b_i \geq b_j$ for all j that is not equal to i. If the bidder still has the highest bid, then that bidder is still winning. That is, because $b'_i > b_j$ for all j that is not equal to i, bidder i will still win. Therefore, the bidder cannot switch from winning to losing while increasing their bid. There are other monotone allocation rules. For example, an allocation rule can assign a value to each bidder that corresponds to their identification number times their bid. In other words, pick the bidder i whose bid b_i maximizes $i*b_i$. This is not a particularly fair allocation rule, but it is still monotone since if a bidder is currently winning, their assigned value must be the highest since their identification number is the same and they are increasing their bid. That is, $i*b_i \geq j*b_j$ for all j that is not equal to i. If they increase b_i to b'_i while keeping all other bids fixed, then their $i*b'_i$ is still the highest (that is, their new $i*b'_i > i*b_i \geq j*b_j$ for all j that is not equal to i). If they have the highest $i*b'_i$, then they are still winning. Therefore, this rule is also monotone. Some allocation rules are not monotone. For example: "If anyone bids more than $10, give the item to Bidder One. Otherwise, give the item to Bidder Two." Then Bidder Two can switch from winning to losing by increasing their bid. Indeed, imagine that all bidders initially submit a bid of $5. Then Bidder Two wins. But if Bidder Two increases their bid from $5 to $15, Bidder One will now win. Because there is an instance where Bidder Two can increase their bid and switch from winning to losing, this allocation rule is not monotone.

[18] The minimum bid to win is also called the "critical bid." Google uses the term "minimum bid to win" in their online documentation.

[19] Note that the minimum bid to win is determined by other bidders' bids and not one's own. It is determined exclusively by the allocation rule of the auction and does not depend on the payment rule. To be more mathematically precise, Bidder i's minimum bid to win $C(b_1,\ldots, b_{i-1}, b_{i+1},\ldots, b_n)$ is a number b such that on bids $b_1,\ldots, b_{i-1}, b_i, b_{i+1},\ldots, b_n$: Bidder i certainly wins when $b_i>b$, certainly loses when $b_i<b$, and may win or lose when $b_i=b$.

[20] One can also consider the effect of ties here. Imagine there are five bidders, who submit bids of $1, $8, $8, $5, $2, respectively. Then all bidders have a minimum bid to win of $8 (for all bidders, the highest other bid is $8). This means that, for example, fixing the bids of Bidders Two through Five at $8, $8, $5, $2, Bidder One will certainly win if they submit a bid strictly larger than $8, and certainly lose if they submit a bid strictly smaller than $8 (depending on how ties are broken, they may or may not win with a bid of exactly $8). Observe that Bidder Two's bid ($8) is exactly equal to their minimum bid to win ($8), and Bidder Three's bid ($8) is also exactly equal to their minimum bid to win ($8). Therefore, Bidders Two and Three may or may not win, depending how ties are broken. All other bidders' bids ($1, $5, $2) fall below their minimum bid to win ($8), and therefore they lose.

22. The concept of minimum bid to win illustrates the main difference between the first- and second-price auction formats: The first-price auction charges the winner their bid, while the second-price auction charges the winner their minimum bid to win.

B. Overview of Reserve Prices

23. A seller can introduce a **reserve price** *r,* to an auction which can be conceptualized in two equivalent ways:[21]

a. The auctioneer enters the auction themselves as Bidder Zero with a bid of *r*, and otherwise the auction proceeds as normal.

b. All bidders whose bid falls below *r* are immediately removed, and all remaining bidders face a "price floor" of *r*. The original single-item auction executes on the remaining bidders. If this results in any bidder winning at a price less than *r*, their payment is increased to *r* (if no one wins the auction or the winner pays more than *r*, then the auction concludes without modification).[22]

24. In a **first price auction with reserve price** *r*, the auctioneer still solicits a bid from each bidder, but the allocation and payment may change. If the highest bid exceeds *r*, then the highest bidder wins and pays their bid to the auctioneer. However, if the highest bid falls below *r*, then the auctioneer keeps the item, and no payments are made. The reason for this can be illustrated by the two interpretations of the reserve price described above:

a. Suppose the auctioneer is referred to as Bidder Zero and each of the other bidders as "real bidders." Bidder Zero, the auctioneer, will submit a bid equal to *r.* If the highest bid from a "real bidder" exceeds *r*, then it is the highest bid overall (*i.e.*, including Bidder Zero's, the auctioneer's, bid which is equal to *r*), so this "real bidder" wins and pay their bid to the auctioneer. If the highest bid from a "real bidder" falls below *r*, then Bidder Zero (the auctioneer) is the highest bidder, so the auctioneer keeps the item, and no payments are made.[23]

[21] The first interpretation is the cleanest to conceptualize. The second interpretation generalizes more naturally to personalized reserves (reserves are defined below), and to a wider range of auction formats. Both interpretations are presented in order to build better understanding.
[22] At the end of this subsection, I provide a conceptually cleaner version of this interpretation. For now, I present the mechanically implementable version of this interpretation.
[23] Alternatively, one can think of this as the auctioneer "wins" the item and keeps it and pays r to themselves.

b. If the highest bid exceeds *r*, then the highest bidder (and perhaps others) is entered into the first-price auction. The highest bid will win and pay its bid (which is larger than *r*, so it does not need to be increased) to the auctioneer. If the highest bid falls below *r*, then no bids survive, and therefore no one gets the item or pays the auctioneer.[24]

25. Similarly, in a **second-price auction with reserve price** *r*, the auctioneer also still solicits a bid from each bidder, but the allocation and payment may change. If the highest bid exceeds *r*, then the highest bidder wins and pays r or the highest other bid, whichever is the greater between the two, referred to as the maximum. If the highest bid falls below *r*, then the auctioneer keeps the item, and no payments are made. Again, the reason for this can be illustrated by the two interpretations of the reserve price described above:

a. Suppose the auctioneer is referred to as Bidder Zero and each of the other bidders as "real bidders." Bidder Zero, the auctioneer, will submit a bid equal to *r*. If the highest bid from a "real bidder" exceeds *r*, then it is the highest bid overall (*i.e.*, including Bidder Zero's, the auctioneer's, bid which is equal to *r*), so this "real bidder" wins and pay the value of the highest other bid among both the other "real bidders" and Bidder Zero. If the highest bid from a "real bidder" falls below *r*, then Bidder Zero (the auctioneer) is the highest bidder, so the auctioneer keeps the item, and no payments are made.

b. If the highest bid exceeds *r*, then the highest bidder (and perhaps others) is entered into the second-price auction.

   i. If the highest other bid falls below *r*, then the highest bidder is alone in the second-price auction, wins with a tentative payment of $0 because there are no other bidders that participate, but then has their payment increased to the price floor of *r*. In this case, the maximum of *r* and the highest other bid is *r*, and the payment is *r*.

   ii. If the highest other bid exceeds *r*, then the highest bidder wins the second-price auction at a price equal to the highest other bid (which already

[24] Observe that a first-price auction with reserve $0 operates the same as the initial definition for the first-price auction. In the first interpretation, Bidder Zero is never the highest bidder with a bid of $0. In the second interpretation, all bids survive, and the winning bid is always already greater than $0.

exceeds *r*, and so does not need to be increased). In this case, the maximum of *r* and the highest other bid is the highest other bid, and the payment is the highest other bid.

iii. If the highest bid falls below *r*, then no bids survive, and therefore no one gets the item or pays the auctioneer.

26. In a first-price auction, there are two relevant ranges for the reserve: (1) the highest bid exceeds the reserve, in which case the auction concludes identically as if there were no reserve or (2) the highest bid falls below the reserve, in which case the auction is essentially nullified, and the item stays with the seller. In a second-price auction, there are three relevant ranges for the reserve: (1) the second-highest bid exceeds the reserve, where the auction concludes identically as if there were no reserve or (2) the second-highest bid falls below the reserve, but the highest bid exceeds the reserve, so the highest bidder still wins, but pays the reserve which is greater than the second-highest bid or (3) the highest bid falls below the reserve, so the auction is essentially nullified, and the item stays with the seller.

27. The leading example can be modified with reserve prices. Imagine that there are five bidders, who submit bids of $1, $8, $3, $5, $2, respectively. In a first-price auction with a reserve price of $4, Bidder Two wins and pays $8. In a second-price auction with reserve of $4, Bidder Two wins and pays $5. That is, both auctions conclude exactly as if there were no reserve, because the reserve is smaller than the second-highest bid. If the auctioneer sets a reserve of $6, then in a first-price auction, Bidder Two wins and pays $8. The first-price auction concludes exactly as if there were no reserve, because the reserve is smaller than the highest bid. In a second-price auction with a reserve of $6, Bidder Two wins and pays $6. That is, Bidder Two still wins because they outbid the reserve. However, Bidder Two pays more because the reserve is treated as the second-highest bid. Observe also that $6 is Bidder Two's minimum bid to win, because $6 exceeds all other bids, Bidder Two will win if and only if they submit a bid above the reserve of $6. If the reserve is $10, then in both the first- and the second-price auction the item remains unsold because the reserve exceeds the highest bid.[25] Figure 2 below illustrates this example.

[25] I further analyze this numeric example in the Appendix C.

Figure 2: The reserve price can act as a clearing price in a second-price auction

	Winner	Payment
First-Price Auction	Bidder Two	$8
Second-Price Auction	Bidder Two	$6

28. The following interpretation of reserve prices may be conceptually more straightforward in some applications. Adding a reserve of r to the first- or second-price auction changes the allocation rule from 'highest bid wins' to 'highest bid at least as large as r wins' (if no such bidder exists, no one wins). However, the winner of a first-price auction with reserve r still pays their bid, and the winner of a second-price auction with reserve r still pays their minimum bid to win.[26]

1) Personalized Reserve Prices

29. A seller can also introduce **personalized reserve prices**, r_i for each bidder i, to an auction. Under personalized reserve prices, all bidders i whose bids falls below r_i are immediately removed. The original single-item auction executes on the remaining bidders. If this results in Bidder i winning at a price less than r_i, their payment is increased to r_i (if no one wins the auction,

[26] More specifically, the allocation rule of a second-price auction with reserve r awards the item to the highest bidder whose bid exceeds r, if one exists, and to no one otherwise. The payment rule of a second-price auction with reserve r charges the winner their minimum bid to win. The allocation rule of a first-price auction with reserve r awards the item to the highest bidder whose bid exceeds r, if one exists, and to no one otherwise. The payment rule of a first-price auction with reserve r charges the winner their bid (if one exists). The numerical example in the Appendix C further establishes the connection between second-price auctions with reserve and minimum bids to win.

or the winner i pays more than r_i, then the auction concludes without modification). The only difference between the identical and personalized reserves is that now the reserve may be different for each bidder.

30. The two most popular auction formats can be augmented with personalized reserves. In a **first-price auction with personalized reserves** ($r_1, \ldots, r_n$), the auctioneer solicits a bid from each bidder. All bidders who do not exceed their personal reserve are immediately removed. The highest remaining bidder (if any) wins the item and pays their bid.[27] In a **second-price auction with personalized reserves** ($r_1 \ldots, r_n$), the auctioneer solicits a bid from each bidder. All bidders who do not exceed their personal reserve are immediately removed. The highest remaining bidder (if any) wins the item. They pay the maximum of their personalized reserve and the highest other remaining bid.[28]

31. The leading example presented above can be modified with personalized reserves as well. Imagine that there are five bidders, who submit bids of $1, $8, $3, $5, $2, respectively. I will consider a few possible settings of personalized reserves. Consider first that the auctioneer sets personalized reserves $1, $10, $1, $8, $1. Then in both the first- and second-price auctions with personalized reserves, Bidders Two and Four are immediately removed as their bids ($8, $5) fall below their personalized reserves ($10, $8). Bidders One,[29] Three, and Five remain. Among them, Bidder Three is the highest bidder. Therefore, Bidder Three wins the first-price auction with personalized reserves and pays $3 (her bid). Bidder Three also wins the second-price auction with personalized reserves and pays a tentative price of $2 (the highest other remaining bid). Because $2 exceeds her personalized reserve of $1, her final payment is $2. One can also compute Bidder Three's minimum bid to win as $2. If she were to submit a bid less than $2, then Bidder Three would no longer be the highest remaining bidder. Similarly, if she submits a bid greater than $2, then she clears her reserve and is the highest remaining bidder.[30] This example is illustrated in Figure 3 below.

[27] Observe that a first-price auction with personalized reserves ($r, r, \ldots, r$) is the same as a first-price auction with reserve r.

[28] Observe that a second-price auction with personalized reserves ($r, r, \ldots, r$) is the same as a second-price auction with
reserve r.

[29] Like with reserves, it is common to break ties in favor of bidders when comparing their bids to their personalized reserves.

[30] We can similarly calculate the minimum bid to win of Bidder One as $3, Bidder Two as $10, Bidder Four as $8, and Bidder Five as $3. For Bidders One and Five, they just need to beat Bidder Three's high bid of $3, because if they do this, they will certainly also clear their personalized reserve. For Bidders Two and Four, they just need to beat their personalized reserves of $10 and $8, because if they do this, they will certainly also be the highest remaining bidder.

Figure 3: The highest other bid determines the clearing price in a second-price auction with personalized reserves

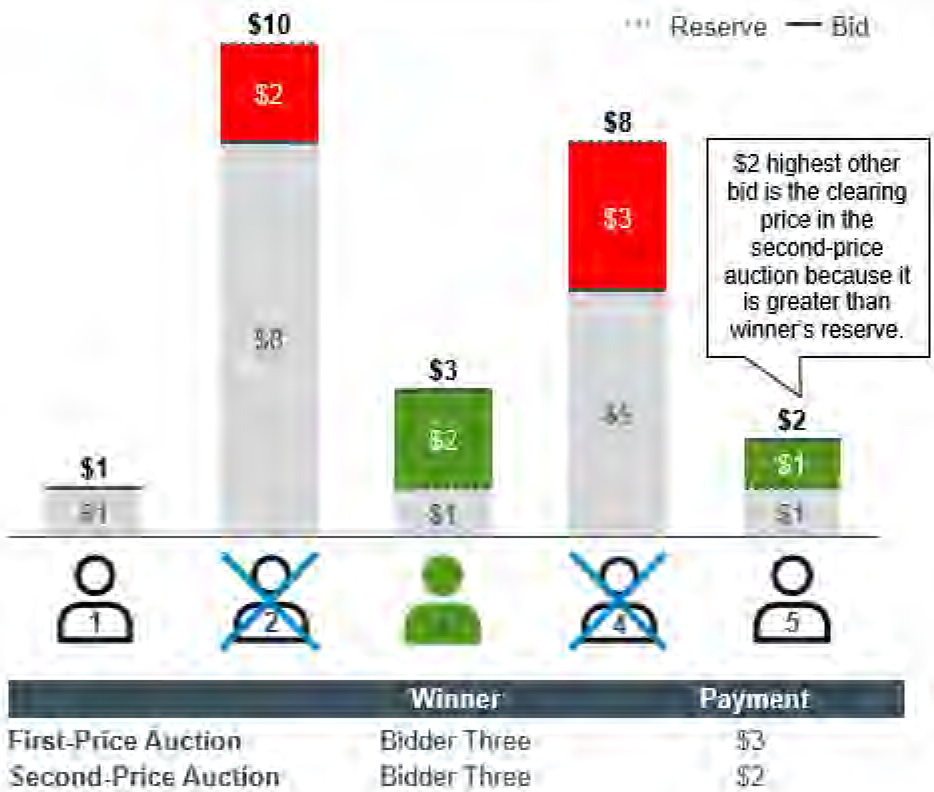

32. Alternatively, consider next that the auctioneer sets personalized reserves $1, $10, $1, $4, $1. Then in both the first- and second-price auctions with personalized reserves, Bidder Two is immediately removed, and Bidders One, Three, Four and Five remain. Among them, Bidder Four is the highest bidder. Therefore, Bidder Four wins the first-price auction with personalized reserves and pays $5 (her bid). Bidder Four also wins the second-price auction with personalized reserves and pays a tentative price of $3 (the highest other remaining bid). Because Bidder Four's personalized reserve is $4, her final payment is increased to $4. We can also compute Bidder Four's minimum bid to win as $4; if she were to submit a bid less than $4, then Bidder Four would not clear her personalized reserve and lose. If she submits a bid greater than $4, then she clears her reserve and is also the highest remaining bidder (because the highest remaining other bidder is $3).[31] This example is illustrated in Figure 4 below.

[31] We can similarly calculate the minimum bid to win of Bidder One as $5, Bidder Two as $10, Bidder Three as $5, and Bidder Five as $5. For Bidders One, Three, and Five, they just need to beat Bidder Four's high bid of $5, because if they do this, they will certainly also clear their personalized reserve. For Bidder Two, they just need to beat their personalized reserve of $10, because if they do this, they will certainly also be the highest remaining bidder.

Figure 4: A personalized reserve price can be the resulting clearing price in a second-price auction

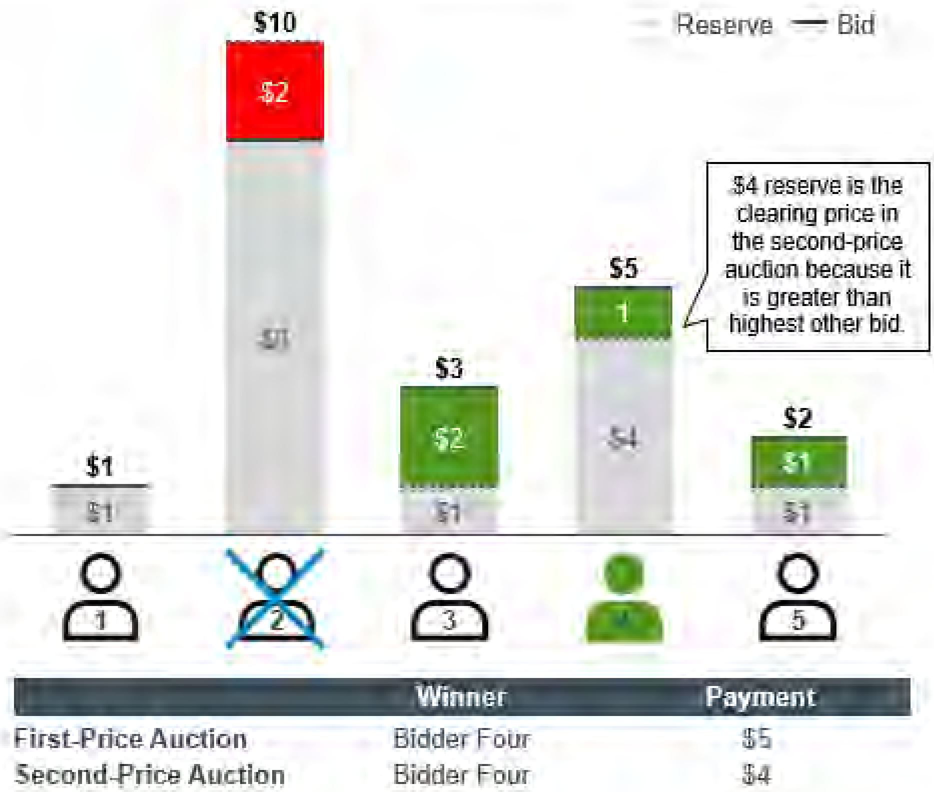

33. Finally, consider that the auctioneer sets personalized reserves $6, $10, $6, $6, $6. Then no one exceeds their personalized reserve, and so the item remains unsold, and no payments occur.[32, 33] This is illustrated in Figure 5 below.

---

[32] We can calculate the minimum bid to win of Bidder One as $6, Bidder Two as $10, Bidder Three as $6, Bidder Four as $6, and Bidder Five as $6. For all bidders, they just need to beat their personalized reserve, because if they do this, they will certainly also be the highest remaining bidder (as there are no other remaining bidders).

[33] The example demonstrates that personalized reserves result in less structure than identical reserves. For example, the first instance shows that with personalized reserves the highest bidder need not win the auction if they bid below their personalized reserve. Indeed, Bidder Two is the highest bidder, yet Bidder Three won the item. The next instance shows that in a second-price auction with personalized reserves, the winner's personal price floor might still bind even when other bidders survive their personalized reserves. Recall that this could not happen with identical reserves. If multiple bidders survived the reserve, then by definition the highest other remaining bid exceeds the reserve and therefore the highest other bid paid by the winning bidder does not need to be increased to meet the reserve. In these cases, the winner pays the reserve only as the only bidder to survive the reserve. The final instance shows that with personalized reserves, the item may remain unsold even if one bidder submits a bid exceeding another personalized reserve.

Figure 5: An auction fails to clear if no bidder passes their personalized reserve

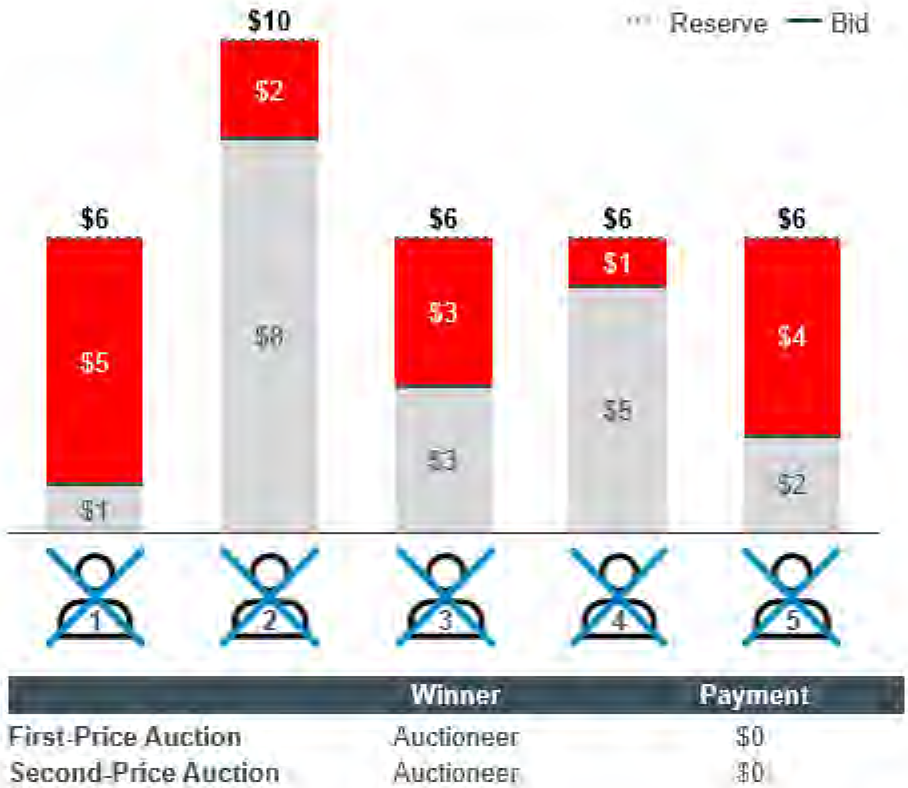

34. There is an alternative interpretation of personalized reserves which can be more straightforward in some applications. Adding personalized reserves of $(r_1,\ldots,r_n)$ to the first- or the second-price auction changes the allocation rule from 'highest bid wins' to 'highest bid at least as large its personalized reserve wins' (if no such bidder exists, no one wins). However, the winner of a first-price auction with reserves $(r_1,\ldots,r_n)$ still pays their bid, and the winner of a second-price auction with reserves $(r_1,\ldots,r_n)$ still pays their minimum bid to win.

35. Table 1 below serves as a reminder of first- and second-price auctions and the impact of reserve prices.

Table 1: Classification of the winner and payment results in auctions with different reserve settings

	First-Price Auction	***Second-Price Auction***
No Reserve	Highest bid wins, **pays bid**.	Highest bid wins, **pays minimum bid to win**.
Reserve r	Highest bid above *r* wins, **pays bid**.	Highest bid above *r* wins, **pays minimum bid to win**.
Personalized reserves ($r_1,\ldots,r_n$)	Highest bid above personalized reserve wins, **pays bid**.	Highest bid above personalized reserve wins, **pays minimum bid to win**.

2) Benefits of Reserve Prices

36. There are multiple reasons why utilizing reserve prices is useful for sellers. Imagine that a homeowner is looking to sell their home. If the seller sets no reserve price, this means that they would accept the best offer received by some deadline. However, the best offer may not be attractive so they may not wish to sell. In particular, the seller may value the home at a particular amount (e.g., at least the amount they currently owe on their mortgage) and they may not wish to sell at a price below that amount. Imagine that the home is worth $50,000 to the seller, based on this the seller might set a reserve of $50,000 to avoid selling their home to a best offer of $10,000. This is better than selling without a reserve at all, but is still suboptimal, since reserve prices are a useful tool to extract extra revenue from sales; in so-called "thin" markets, reserve prices are essential (imagine that only a single potential home buyer shows up, the seller would be better off setting a reserve to get a reasonable offer). Generally speaking, thin markets are markets where the number of competitive buyers is relatively low. Note that a market with many potential buyers can still be thin if a small number of them have significantly higher value than the rest. In so-called "thick" markets, reserve prices are less essential but still useful (imagine that hundreds of potential home buyers show up, there is likely plenty of competition for high bids already even without a reserve, but a reserve is still useful in the unlucky event that only a few of the buyers turn out to be serious). Generally speaking, thick markets are markets where the number of competitive buyers is relatively large.

37. Still, using reserves effectively always requires careful optimization.[34] For example, if potential homebuyers see the seller setting a low reserve in a first-price auction, they will submit low-ball offers, and the seller will make less money selling their home even if bidders would have been willing to pay a lot more. Setting reserves is a tricky business; set the reserve too low and the seller misses out on extra revenue but set the reserve too high and the seller will miss the sale entirely.

38. Reserves play different but similar roles in first- and second-price auctions. In a second-price auction, the role of the reserve is to increase revenue in case it falls between the highest and the second highest bids. But this comes with the possibility of setting a reserve that is higher than the top bid and failing to sell the item, which would decrease the seller revenue to $0. In particular, the bidders do not even need to know the reserves for reserves to be revenue-relevant for the seller in a second-price auction, since the reserve price directly determines the sale revenue when it is in between the highest and the second highest bids. In a first price auction, the reserves are useful due to their effect on the bidders' behavior.[35] In particular, if the reserves are known to the bidders beforehand, the reserves impact the submitted bids. For example, in the home selling context, the seller would be better off revealing the minimum price they are willing to accept to the bidders, since it will affect the submitted bids for the house. Like in a second-price auction, a too-high reserve risks the item going unsold in a first-price auction as well.

39. Similarly, there are multiple reasons why utilizing personalized reserve prices is useful to sellers. Imagine the pricing strategy of an airline company. The company may wish to charge more to a buyer traveling for business (with the expectation that the traveler's company will pay the bill) versus traveling for pleasure (expecting that the traveler will pay the bill themselves), to a buyer booking at the last minute (in desperation for a flight to an important event) versus far in advance (likely with numerous options), or to a buyer who contacted the seller quickly through the website (who did not seem to look particularly hard for a deal) versus a buyer who contacted the seller through a convoluted discount deals site (who presumably looked quite hard to find that deal). These cases are not fully explained by a preference for one buyer over the other but are better explained by the fact that some buyers are likely willing to pay more than other buyers, and the seller could therefore extract more revenue by setting a higher price to them.

---

[34] Myerson's seminal work (1981) describes how to set the optimal reserve in several settings, based on aggregate market data about potential bidders' values. Roger B. Myerson. "Optimal Auction Design." *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73.

[35] As I explain below, this is not the case in a second-price auction, since a bidder's best action is bidding their value for the auctioned item, no matter the reserve.

40. This reasoning carries over to personalizing reserves in auctions as well. Imagine the seller is auctioning a weekend in their apartment to two bidders, one of whom is a partier and the other of whom is quiet, using a second-price auction. A reasonable strategy might be to set a very high reserve for the partier, and a smaller reserve for the quiet one (the mindset being that only if the partier is willing to pay a lot, then the seller is happy to have them, and otherwise would prefer the quiet one). A similar strategy would make sense for business versus pleasure travelers because the seller would want the opportunity to set a high personalized reserve for the likely price-insensitive business traveler but set a comparatively lower personalized reserve for the leisure traveler (in case the reserve overshoots the business traveler's valuation). The example shows that personalized reserves make sense both to express preferences for one buyer over another (*i.e.,* if one buyer might cause physical damage to the apartment and the other might not), and to optimize the seller's revenue even without preferences for one buyer over another.

C. Strategic Behavior in Auctions

41. Incentives that different auction formats foster in bidders are an important differentiator among auction formats. In order to understand the strategic implications of auctions, a measurement of how the agents evaluate the outcomes of auctions is needed. A bidder's **value** or **willingness to pay** for an item is a number v such that:

a. For any price p larger than v, if given the choice to win the item and pay p or lose, the bidder prefers to lose.

b. For any price p smaller than v, if given the choice to win the item and pay p or lose, the bidder prefers to win the item and pay p.

c. If given the choice to win the item and pay a price that is equal to v or lose, the bidder is indifferent.

42. The concept of a willingness to pay is intuitive, although perhaps not with this level of precision. For example, any time someone decides that they are willing to pay $5 for a sandwich, they are implicitly deciding that their value for a sandwich exceeds $5. Any time they decide they are unwilling to pay $10 for a sandwich, they are implicitly deciding that their value for a sandwich is less than $10. If they spend further time deciding, they might narrow the range further, but are unlikely to decide on the exact value. In reality, for individuals it is not worth the effort to nail down precise values on everyday purchases. Likewise, big purchases made by individuals rarely turn out so that the price faced is extremely close to the individual's value, so they may not go through

the process of determining precise value for the good. However, for businesses that regularly participate in millions of small-scale auctions, determining a precise value for a good is very important. For example, if the item for sale is the right to display an ad to an online user, a business likely has data analytics predicting the downstream profit it expects to earn from this ad and is happy to purchase the display rights if and only if the predicted profit exceeds the price.[36]

43. There are two canonical models for how bidders form their values. First, bidders have **private values** in a single-item auction if each bidder has all the information necessary to determine their value for the item.[37] As a special case of private values, if it is further the case that knowing one bidder's value for the item provides no meaningful information to discern another bidder's value, the bidders have **independent private values**.[38] Second, bidders have **interdependent values** in a single-item auction if one bidder has information that informs another bidder's value for the item.[39] In the interdependent values model, each bidder has a private "signal" known only to them which contains information about the value of the item that is being auctioned. Each bidder then forms their value for the item as a function of all bidders' signals. Private values correspond to the special case where a bidder's valuation of the item ignores all other bidders' signals and depends only on that bidder's signal.[40]

44. To illustrate the nature of independent private values, imagine an apple seller. Every customer who walks into the store has a different willingness to pay for apples; some customers love apples and are super hungry while others could not be paid to eat one. But every customer is perfectly capable of forming their own value for an apple based entirely on their own taste and perception of outside options, so apples are an instance of private values. Moreover, when two (unrelated) customers walk into the apple store, imagine the seller learns that one has a value of

[36] There are of course complications that prevent the ad buying story from being quite so simple. Some complications will be discussed later, with the context of the case. For now, I simply want to make the point that the concept of a value is well-defined and relevant.

[37] This model is typically attributed to seminal work of Vickrey (1961) that significantly contributed to his 1996 Nobel Memorial Prize in Economic Sciences. William Vickrey. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *The Journal of Finance* vol. 16, no. 1. 1961. pg. 8-37.

[38] This model is typically attributed to seminal work of Myerson (1981) that significantly contributed to his 2007 Nobel Memorial Prize in Economic Sciences. The precise mathematical meaning of this will be discussed later in the detailed analysis. Roger B. Myerson. "Optimal Auction Design." *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73.

[39] This model is typically attributed to seminal works of Wilson (1969), and Milgrom and Weber (1982) that significantly contributed to Milgrom and Wilson's 2020 Nobel Memorial Prize in Economic Sciences. Robert B. Wilson. "Competitive Bidding with Disparate Information." *Management Science* vol. 15, no. 7. 1969. pg. 446–448. Paul R. Milgrom and Robert J. Weber. "A Theory of Auctions and Competitive Bidding." *Econometrica* vol. 50, no. 5. 1982. pg. 1089–1122.

[40] In the context of the auctions for online display ads, I believe it is best to conduct the analysis assuming that advertisers have independent private values for impressions. The only exception is when I consider the potential impact of Enhanced Dynamic Allocation on direct deals in Section IV.

$2.28 for an apple. This information does not say much about the other customer's valuation, since the seller already has aggregate data on potential customers' apple values, and the fact that one customer happens to be on the high-end of that dataset tells nothing about where an unrelated customer might lie. So, apples are an instance of independent private values. To illustrate the nature of (non-independent) private values, consider instead a person who wound up with a friend's used car and is hoping to sell it to one of several used car experts who can perfectly determine its quality and resale value. In this case, each expert still has a private value for the car since they gain no information from other experts, because their own assessment is perfectly sufficient. However, the private values are not independent because the seller initially has no idea whether the used car is a "lemon," and therefore *a priori* has uncertainty over how much an expert might be willing to pay. If the seller learns that one expert has a high value for the car, they will conclude that it is not in fact a lemon and expect other experts to have high values as well. Or it might be the case that the experts cannot make a perfect assessment, they can only perform an inspection and catch some issues. In this case, the experts' values are interdependent because even after performing their own inspection, one expert's value will go up after learning that another expert found no issues, or down after learning that another expert found significant issues.

45. These examples identify three levels of complexity: Independent private values (selling apples) are the simplest; every bidder knows their own value, and learning one bidder's value doesn't help price apples better for others. Private, but not independent, values are intermediate; every bidder still knows their own value, but the auctioneer can perhaps leverage any knowledge gained from one car expert to better price the item towards others.[41] Interdependent values are complex; even the bidders do not know their own values, and therefore the car experts must think carefully before bidding about how their value will change after learning others' bids at the auction's conclusion.

46. The subject of how bidders bid in auctions is an active research area.[42] However, the key concepts needed at this point are simply that (a) bidders prefer some outcomes over others since bidders prefer to win at a price below their value rather than lose, and prefer to lose rather than

[41] In fact, seminal work of Cremer and McLean (1988) designs a complex auction where bidders place elaborate bets on other bidders' bids that generically outperforms anything simple. Jacques Crémer and Richard P. McLean. "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica* vol. 56, no. 6. 1988. pg. 1247–1257.

[42] For example, Li (2017) and the 599 (at time of writing) subsequent papers. Shengwu Li. "Obviously Strategy-Proof Mechanisms." *American Economic Review* vol. 107, no. 11. 2017. pg. 3257-3287.

win at a price above their value (beyond this, subject to winning, bidders prefer to pay as low a price as possible) and (b) bidders will strategize while bidding in attempt to get preferred outcomes. The complexity of strategies depends on how their values are formed (*i.e.,* private versus interdependent), and the auction format.

47. In order to understand how auction formats affect bidder strategies, further terminology is needed. A sealed bid single-item auction is **truthful** if each bidder receives the best possible outcome (given the other bidders' bids) by submitting a bid equal to their own value.[43]

48. I now provide an example to illustrate the concept of truthfulness. Imagine a single-item auction with two bidders. Bidder One has value $5 and Bidder Two has value $10. Imagine further that the auctioneer has chosen a first-price auction, and Bidder One has submitted a bid of $3. What is the best possible outcome for Bidder Two in such a setting? It may be tempting to first claim that the best possible outcome for Bidder Two is to win the item and pay $0. However, there is nothing Bidder Two can do to make this happen since Bidder One has submitted a bid of $3, so the only outcomes available to Bidder Two are to lose the item (by submitting a bid less than $3), or to win the item at a price greater than $3 (by submitting a bid $b that is higher than $3).[44] In particular, winning the item and paying $10 is certainly not the best possible outcome, which is the resulting outcome should Bidder Two bid their value (a better outcome, for example, would be win and pay $3.01 by submitting a bid of $3.01). Therefore, the first-price auction is not truthful.[45] The incentive of a bidder to submit a bid that is lower than their value for the item is called "**bid shading**." The amount by which the bidders are incentivized to shade their bids is determined by their comparison of the risk of paying more when they submit a higher bid and the risk of losing the auction when they submit a lower bid. Figure 6 below illustrates the ideas in this paragraph.

[43] More formally, Bidder i cannot control what bids are submitted by the other bidders. But, no matter what those bids are, Bidder i can go through the thought process of "which bid b_i gets me the best possible outcome, given that the other bidders have bid $b_1,\ldots, b_{i-1}, b_{i+1},\ldots, b_n$?" An auction is truthful if the answer is always "submitting b_i equal to my value gets me the best possible outcome."
[44] Depending on how ties are broken, perhaps it is possible to win the item at exactly $3.
[45] In particular, note that an auction is truthful if it is always best to submit a bid equal to the value. Because this example witnesses one case where it is not, this example in fact constitutes a proof that the first-price auction is not truthful.

Figure 6: In a first-price auction, the bidders are incentivized to submit bids lower than their values

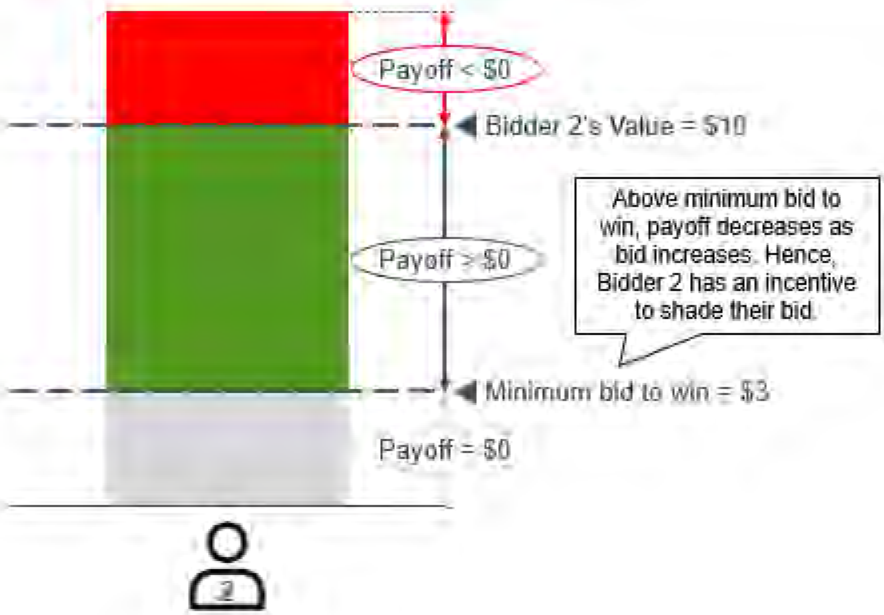

49. Imagine instead that the auctioneer has chosen a second-price auction, and Bidder One has submitted a bid of $3. In this case, the options available to Bidder Two are (a) lose (by submitting a bid less than $3), or (b) win and pay $3 (by submitting a bid greater than $3). Because Bidder Two's value is $10, their preferred outcome of the two possibilities is to win and pay $3. Submitting a bid of $10 is one way for Bidder Two to win and pay $3, and therefore bidding their value is one way for Bidder Two to get the best possible outcome.[46] The implications of these examples are not anomalies. Second-price auctions are truthful, and this is one of their key advantages in comparison to first-price auctions.[47, 48, 49]

[46] Still, recall that an auction is truthful if it is always best to submit a bid equal to the value. This example witnesses one case where it is, and so does not constitute a proof that the second-price auction is truthful. However, it does give intuition for why this is the case.
[47] A complete proof of this claim can be found in Theorem 1 the Appendix C.
[48] The truthfulness of the second-price auctions also holds with interdependent values, as long as the interdependency satisfies a technical condition called "single-crossing". Intuitively, the single-crossing condition assumes that Bidder i's value is more sensitive than Bidder j's value to Bidder i's signal (for example, car expert i places more emphasis on car expert i's inspection than car expert j places on car expert i's inspection). One precise mathematical statement is that for all signals $s_1,\ldots, s_n$, we have $\partial v_i(s_1,\ldots,s_n)/ \partial s_i \geq \partial v_j (s_1,\ldots,s_n)/ \partial s_i$. There are nuances to this generalization beyond the technical condition, it requires the auction designer to know how bidders map each other's signals to values and requires the auction designer to directly solicit signals rather than bids.
[49] Single-item auctions that charge minimum bids to win satisfy an even stronger property that bidding one's value is a dominant strategy, meaning there is a mathematically formal sense in which bidding the value is strictly better than any other bid. Formally, this just means that for any bid b not equal to a bidder's value: (a) bidding their value always gets the bidder an outcome at least as good as bidding b, no matter the other bids, and (b) there are some possible bids of the others where the bidder gets a strictly better outcome bidding their value than bidding b. The precise truthfulness properties that second-price auctions possess will not play a role in later analysis, this is stated primarily

1) Comparison of Incentives in First- and Second-Price Auctions

50. First- and second-price auctions differ in terms of the strategic behavior they elicit in the auction participants and the auctioneer. The suitable format depends on the context of the auction.

51. Given its truthfulness property, the second-price auction format may seem to be the ideal auction format compared to the first-price auctions. This is because bidding in auctions requires strategic sophistication. Bidders do not know other bidders' values, and so they do not know what bid will win them the item at the lowest possible price, in other words, their minimum bid to win. In a first-price auction, bidders would ideally want to bid a penny above their minimum bid to win (if their value exceeds that). However, in a truthful auction, the strategic sophistication is not needed because it is in the bidders' best interest to bid their own valuations of the auction item truthfully. Also, bidders tend to prefer straightforward auctions where they do not need to do any strategic bidding. This further enhances the benefit of the truthfulness of the second-price auction format.

52. However, there are potential mitigating factors as well. Truthfulness only holds when viewing this auction in isolation, it does not necessarily hold when considering a series of auctions for several reasons. One example to have in mind is when the auctioneer might change their reserve in later auctions based on bids in earlier auctions. Once the auctioneer's reserve is fixed, a second-price auction with that reserve is truthful, and a bidder's optimal bid for this auction is their value. However, if the bidder has a high value for the item and bids as such, this may indicate to the auctioneer that they should set a higher reserve in future auctions, and this higher reserve will certainly hurt the bidder in the future.[50] It may also signal to other bidders that similar items are valuable and increase their future values (and therefore bids).[51]

to note that the argument that bidders should bid their value in a truthful auction is slightly stronger than what is implied by Theorem 1 in the Appendix C.

[50] I believe this concept is relevant when analyzing Google's conduct called "Reserve Price Optimization." This is discussed in detail in Section IX.

[51] Additionally, Myerson's (1979) seminal work introduces the so-called "revelation principle", which essentially provides a truthful wrapper to put around a non-truthful auction to handle the non-truthfulness on behalf of the bidder. Essentially, imagine that the real bidder hires a sophisticated data-rich consultant to bid in a non-truthful auction on their behalf, tells the consultant their value, and lets the consultant take care of optimizing the bid. From the bidder's perspective, assuming the consultant is honest and effective, this is now just a truthful auction (the bidder is best off giving the consultant the most accurate information to work with and letting them do their job), and only the consultant deals with the messy optimization. In settings where such a sophisticated data-rich consultant is widely available, the additional benefit of truthfulness is smaller. Roger B. Myerson. "Incentive Compatibility and the Bargaining Problem." *Econometrica* vol. 47, no. 1. 1979. pg. 61–73.

53. Another important point is that first-price auctions are **credible**,[52] but second-price auctions are not. Taken for granted during discussion above is that bidders trust the auctioneer to run whatever auction the auctioneer promises. For example, imagine a scenario where a bidder is participating in (ostensibly) a second-price auction, submits a bid of $20, and then is told that the second-highest bid was $19.99. It is possible that this bidder was participating in a second-price auction with another bidder who submitted a bid of $19.99, or the auctioneer could be running a rigged auction where they submit a shill bid a penny below the highest bidder. There would not be any way to know which one was the case. Informally, this establishes that a second-price auction is not credible. On the other hand, first-price auctions are credible (Akbarpour and Li, 2020).[53] When a bidder submits a bid of $20, they know that they will either win and pay $20 or lose. Moreover, even if a cheating auctioneer wants to run any auction that picks one of these two outcomes for each bidder, their revenue-maximizing option is indeed to just give the item to the highest bidder and charge them their bid.

54. These are two examples of properties that auction theorists consider when evaluating one auction format over another. There are several others, too. For example, although second-price auctions are not difficult to explain, first-price auctions are certainly simpler. Furthermore, bidders may find it convenient that after submitting a bid to a first-price auction, they know that if they win, they will pay exactly that bid (whereas in a second-price auction, the precise payment made is ultimately determined by other bidders). This might be convenient from the perspective of a bidder managing a budget across several auctions, or writing contracts to profit-share with an entity that bids on their behalf.

D. Information in Auctions

55. The amount of information participants have in auctions is an important consideration in the design and analysis of auctions in several aspects, as explained in the paragraphs below. For example, buyers need information to determine their value for the item(s) being sold. Imagine that the item for sale is the right to display a running shoes ad to a particular user. Information about this user, such as whether they like running or not, their age, their spending habits, and their willingness to pay for a pair of running shoes would all be factors in determining the value of this user's attention. Furthermore, buyers need information to determine their optimal bidding strategy

[52] The concept of "credible auctions" is after Akbarpour and Li (2020). Mohammad Akbarpour and Shengwu Li. "Credible Auctions: A Trilemma." *Econometrica* vol. 88, no. 2. 2020. pg. 425-467.

[53] Mohammad Akbarpour and Shengwu Li. "Credible Auctions: A Trilemma." *Econometrica* vol. 88, no. 2. 2020. pg. 425-467.

in a non-truthful auction. If this is a first-price auction, it would surely be great for this bidder to know the other bidders' bids (or any piece of information regarding the bids of the other bidders), so that they do not bid more than what is necessary to win. Lastly, sellers need information to set optimal reserves (or in general, to optimize their revenue). For example, imagine a person who moved to a completely new country and opened a pizza joint. For them, it would be useful to know how much locals typically pay for meals, how enthusiastic they are about pizza, and if they view pizza as an exciting meal or something cheap.

56. Information is useful for bidders when they are determining their value for the auctioned item. If one were to ask a person for their value for an apple, they might be able to narrow it down within one dollar range. If instead they were asked their value for the right to display an ad for their running shoes to an internet user, this is a much harder value to pin down. If they derive value primarily from whether the user chooses to buy a pair of shoes from them, they need to predict the likelihood that this will happen, which depends on if the user likes running, if the user buys new shoes frequently, and if the user tends to click on ads. More granular information would also be useful, such as miles per week this user runs, the last time they bought a new pair of running shoes and the types of ads they previously clicked. There is further information that could continually refine the prediction for the likelihood that the user makes the purchase decision.

57. For more accurate analysis, sources of information should be separated into external from the auction and internal to the auction. An external source of information could be data from past auctions, market research, etc. An internal source of information would be anything revealed by another bidder's actions in this auction itself. Imagine that a bidder participates in an auction for a car. The bidder's first step would be to ask for information about the car, the production year, mileage, etc. They might further ask for an inspection, a service history, etc. All of this is external information which is learned without participating in an auction. Once the auction takes place, if the bidder notices that the other bidders submitted very low bids, this would suggest that they may have done their own inspection and caught a problem with the car, which suggests that perhaps the bidder should lower their own value for the car. This is internal information since it is learned only after observing other bidders' behavior in the auction. This example is also related to the discussion of private and interdependent values. Private values are derived entirely from external information, and interdependent values are derived at least partially by internal information.

58. Information also helps bidders to optimize their bids. One benefit of a truthful auction is that the only "hard part" of the auction is determining the value. Once the value is determined, it is optimal to bid that value. But many auctions are non-truthful. In a non-truthful auction, optimal bidding is challenging since, by definition, it requires knowing competitors' bids. For example, imagine a bidder in a first-price auction, and their value is $5. What should they bid? They definitely should not bid anything above $5. But other than that, any strategy for this bidder depends on other bidders' bids. But, if somehow the others' bids can be seen, then bid optimization becomes a straightforward process (either bidding a penny more than the maximum bid, or purposefully losing if the maximum bid exceeds the willingness to pay).[54] Hence, the ability to see others' bids, especially if the other bidders cannot see rival bids, would be a great strategic advantage.[55] Similarly, there would be benefit from any data that helps predict other bidders' bids (perhaps, as an example, their historical bids submitted on similar items).

59. Lastly, the sellers use information when they are determining appropriate reserve prices for their items. For example, in the context of a second-price auction, a too high reserve will nullify the entire auction, a too low reserve has no impact, and setting a reserve in the sweet spot between the highest and second-highest bid yields extra revenue. For example, if the highest bid is $20, and the second-highest bid is $15, the seller would ideally like to set a reserve at exactly $20. Failing that, they would really like to avoid setting a reserve above $20 and prefer to set larger reserves between $15 and $20. But if instead the highest bid was $100, and the second-highest bid was $50, the seller would ideally like to set a reserve exactly at $100, definitely not over $100, and somewhere between $50 and $100. Notice that a good reserve in the first case ($20) is useless in the second, while a good reserve (greater than $50) in the second case nullifies all revenue in the first. *A priori*, with no further information, deciding on the optimal reserve is challenging at the very least. But every time a seller sees bids in an auction for a similar item, they learn a little bit about what they might expect the next time. This data is valuable, because it allows the seller to predict whether they are more likely to be shooting for a reserve between $15 and $20 or between $50 and $100, and to target their reserve at the likely case.

60. Another crucial concept is **information asymmetry**, which refers to cases where one agent has more information than another. One source of information asymmetry might be if one special bidder gets to see others' bids in a first-price auction before submitting their own. The

[54] In this example, if the highest other bid is $3.38, they should bid $3.39. If the highest other bid is $1.28, they should bid $1.29.
[55] This point is relevant for "Last Look" combined with "Header Bidding" in Google's Ad Server DFP. These conducts will be discussed in detail later.

special bidder is better able to optimize their own bid, and also may be more informed about their own value if values are interdependent. Another example might be if one seller gets feedback from lots of auctions, but another does not. Then, the first seller is better informed and can potentially set better reserves. Imagine someone, while cleaning out their basement, finds their old collection of Pokémon cards and decides to sell them. They would like to set a reserve to make sure they optimize their revenue, but they might have no idea for how much Pokémon cards go these days. Any information gathered will be helpful to stop them from setting a reserve of $100 and risk getting fleeced, or a reserve of $100,000 and risk their collection going unsold. After setting a good reserve, they might launch a first-price auction. A friend of the seller might reach out and express their interest, but also the friend might be unable to both determine what their true value is (part of their value derives from their ability to further resell the set when they lose interest) and to bid strategically once a value is determined. In turn, the seller might decide that they will just share bids with the friend as they roll in and let them submit theirs at the end. This allows the friend to both (a) accurately form their value for the set, by observing the bids of others, and (b) bid optimally, by submitting a bid just a penny above the highest other bid (if they decide they want to win). By sharing this information with their friend, the seller has created an information asymmetry since the friend is both better informed about their value than other bidders, and better able to optimize their bidding strategy.

III. ONLINE DISPLAY ADVERTISING

61. In this section, I provide an overview of parties and concepts that arise within online display advertising.

A. Pertinent Products and Parties in Transactions for Online Display Ads

62. There are several players involved in online display advertising, the goal of this section is to outline those players as well as their goals and incentives.

1) Publishers and Ad Servers

63. **Publishers** are entities (e.g. the New York Times) that have webpages that can display ads and are therefore sellers of inventory, which is effectively the “eyeballs” of their users. Publishers (and third parties) have a range of information about the particular users visiting their page, and each unique visitor can be thought of as its own item for sale. The item for sale is referred to as an **impression**.

64. Like any other business selling goods, in the context of online display advertising, a publisher's primary goal is to make as much revenue as possible by selling their inventory. However, publishers must be mindful that their true goal is long-term revenue, and that sometimes actions that increase immediate revenue may be harmful to long-term revenue. Hence, publishers have concerns such as the quality of ads displayed in addition to immediate revenue. For example, if users find display ads offensive or annoying, they may negatively impact the publisher's brand and future revenue streams.

65. There are three key elements of the display ad sale process from the publisher's side. First, the primary benefit of the display ad ecosystem is that ads are targeted, which means that each user's eyeballs are treated as a truly distinct item. For example, the right to display an ad for running shoes to a runner is more valuable than showing that ad to a golfer. Second, the decision regarding which ad to display happens nearly instantaneously. Once the publisher learns of a user visiting its website, it must decide which ad to display by the time the webpage loads. Third, the publisher may not necessarily have a natural network of advertisers ready and waiting to bid on its impressions. The publisher must somehow reach interested advertisers.

66. Publishers face a challenging task when they are trying to sell impressions. Imagine the selling process of a publisher, who needs to track, store, and share targeting data on each user visiting their website, reach a wide network of potential advertisers, decide a revenue-maximizing auction to run, and do it all nearly instantaneously while the webpage loads. Publishers usually outsource these tasks to a dedicated product called an ad server.

67. An **ad server** is a service that helps publishers manage and sell inventory, overcoming the technical challenges listed above. Inventory management and optimal pricing are also data-intensive (to accurately determine the aggregate market for each impression) and mathematically sophisticated (to determine what the optimal auction is given the data).

68. An ad server's revenue typically comes either by charging publishers a fixed monthly rate, or charging a fee based on the volume of impressions served.[56] This suggests, in principle, that an ad server's primary goal is to serve their publisher's goals, since the ad server's monetary

---

[56] *See* AdGlare. "Plans & Pricing." Accessed on May 31, 2024. https://web.archive.org/web/20231203094651/https:/www.adglare.com/pricing (describing that AdGlare charges a monthly rate that is based on the number of ad requests.)

incentives are aligned with the publisher's monetary incentives.[57] Google operates an ad server called DoubleClick for Publishers (DFP), which was later merged into Google Ad Manager.[58]

69. There are two key ways through which an ad server might sell inventory. In a **live ad auction**, the ad server learns that a particular user is visiting the publisher's webpage, and while the page loads, runs an auction for the right to display an ad to this particular user. The auction begins only after the user is known, so potential advertisers can submit a bid based on the fine-grained information they learn about the user. In a **direct deal**, the publisher pre-arranges a contract with an advertiser to display their ad some number of times across some period at some predetermined price per impression, perhaps to users that satisfy some coarse targeting criteria. The ad server manages that deal as users visit the publisher's webpage. Because of this, while targeting criteria can still be used, it is coarser in comparison to the real-time data available in a live ad auction.[59, 60]

70. Establishing a relationship with advertisers for direct deals is time consuming and requires a business network. It is perhaps worth the effort for large publishers such as the New York Times and large advertisers such as Nike, but it is unlikely that smaller publishers (for example, a local food blog) grab the attention of Nike's marketing department for direct deals or smaller advertisers (for example, a local escape room) grab the attention of The New York Times for direct deals. Because live ad auctions are automated, there is no reason why code written by a local food blog cannot interact with code written by Nike.[61] As a result, live ad auctions enable small publishers

[57] An ad server that is not standalone (i.e., owned by an entity that operates elsewhere in the online display ad auction ecosystem) could certainly have alternate primary goals. These could be benign and still aligned with a publisher's if that entity views ad servers as a necessary part of the ecosystem for their primary service to thrive. But this also raises potential for misaligned incentives.

[58] Along with Google's own ad exchange, in 2018. This will be explained further below. *See* Jonathan Bellack. "Introducing Google Ad Manager" (June 27, 2018). Accessed on May 31, 2024. https://web.archive.org/web/20240112234145/https:/blog.google/products/admanager/introducing-google-ad-manager/

[59] NT Technology. "Why Is Targeting in Programmatic Ads Better Than Usual?" Accessed on June 5, 2024. https://web.archive.org/web/20240228203217/https://nt.technology/en/faq-en/why-is-targeting-in-programmatic-ads-better-than-usual/ ("Improved targeting capabilities. Through programmatic advertising, you can easily target the specific audience you want to reach using all target opportunities. Rather than trying to reach sports car fans on an auto site, brands have the opportunity to create an audience segment of sports car fans and reach them across hundreds of websites, wherever they happen to be online.")

[60] In reality, there are more types of trade in online display advertising markets, but these two are the most relevant ones to the case, so I choose to focus on these. *See generally* Google. "Line item types and priorities." Accessed on May 31, 2024. https://web.archive.org/web/20240216154938/https:/support.google.com/admanager/answer/177279?hl=en

[61] But even this is still not trivial, since something still has to do the work of finding and connecting these two codebases together (and quickly, by the time a webpage loads). This challenge motivates the role of exchanges, which I discuss after introducing the buy side of the market.

to display ads from a wide range of advertisers without investing in the business network aspect of advertising.

71. A publisher can consider both a direct deal and a live ad auction for the same impression. For example, upon learning that a user is visiting their webpage, a publisher could first check if the impression satisfies a high-value direct deal and if so, sell it via direct deal. If not, the publisher could sell the same impression via live auction instead.

2) Advertisers, Large and Small Ad Buying Tools

72. **Advertisers** are entities that wish to display ads to users and are therefore buyers of inventory. Each advertiser has a distinct value[62] for each impression and that value is determined based on the information the advertiser learns about the user behind that impression. An advertiser's primary goal is to win impressions at a price below their value and to pay as low a price as possible.

73. Advertisers can purchase impressions through direct deals or live ad auctions, which can be challenging tasks. Imagine that Nike is looking to purchase targeted online advertising. They might engage in traditional marketing with large publishers via direct deals, but they will need to engage with live ad auctions to reach small and midsized publishers. Nike needs to process data on each impression to determine its value, manage an advertising budget across an extended time horizon, optimize bidding strategies in live ad auctions, and get connected to publishers in the first place. As a result, advertisers outsource these tasks to dedicated products called ad buying tools.

74. An **ad buying tool** is a service that helps advertisers find impressions that are available for sale, bid appropriately to balance the likelihood of winning versus price paid, manage a budget across a time horizon, process any available data on the impression to inform their value, and generally manage the process of purchasing impressions via live ad auctions. An ad buying tool's

[62] In the terminology introduced in Section II, I assume for the majority of this report that the advertisers have independent private values for impressions (the only exception is when identifying a potential impact of Enhanced Dynamic Allocation on direct deals in Section IV). This is a simplifying assumption (it is likely that no auction in real life purely abides by the independent private values model) that makes the analysis more tractable, and it is a sensible assumption to make since (a) internal Google documents demonstrate that Google assumes this as well (*e.g.*, GOOG-AT-MDL-004016180), (b) bidders are heterogeneous in their valuation for impressions (the impression from an avid runner is valued differently by Nike and McDonald's) and since the bidders do not know each other's identities, even if they learn about others' bids it could possibly not be that useful towards determining their own valuation. GOOG-AT-MDL-004016180 at -94. February 20, 2020. "Auction Theory Primer." ("Unlike with 2nd-price, in order to bid [in a first-price auction], *buyers must believe something about the competition!* **Assumption:** [emphasis in original] Independent Private Values model.")

primary goal is to optimize its revenue, which is typically earned as a fraction of payments made by the advertisers it serves (see the examples in paragraph 80 for a numerical illustration). Google has ad buying tools in two markets. Google Ads serves small advertisers and DV360 serves large advertisers.

75. Arguably the most important function of an ad buying tool is that it determines the bids on behalf of the advertisers, according to the goals they input to the system.[63] Since online ad auctions happen almost instantaneously, the advertisers themselves cannot possibly submit bids into the auctions. Instead, they input their goals for their advertisement campaigns into the ad buying tool, which then comes up with bids in a timely manner when impressions become available. The advertiser goals usually include parameters like the desired volume of impressions, budget and time horizon allocated for the campaign, and targeting criteria.[64]

3) Ad Exchanges

76. Publishers (via ad servers) and advertisers (via ad buying tools) form the sell side and buy side of the markets for live ad auctions. It is not a trivial process for ad servers and ad buying tools to find each other and transact. Even for something that is commonly bought and sold, such as a designer coat, finding every interested buyer on the internet is a difficult task. As a buyer, it is also a difficult task to scour the internet to find all the designer coats you are interested in. Hence a third-party market/exchange/bazaar would be relied on to aggregate supply and demand. For example, customers go to platforms like eBay for Pokémon cards, Etsy for engraved chopsticks, and Amazon for books. In all of these cases, the customers rely on the platform primarily to match them to sellers.[65] The market for impressions is no different, and ad exchanges exist to help publishers meet advertisers.

77. **Ad exchanges** provide the service of matching advertisers (buyers) to publishers (sellers). Ad servers contact an ad exchange with inventory for sale, and the exchange then connects to ad buying tools. Importantly, ad exchanges do not merely connect advertisers to publishers, they also run an auction to determine which advertiser wins the impression. That is, an ad exchange is more like the New York Stock Exchange (which specifies the stock trading mechanism) or Uber and Lyft (which specify the price at which riders and drivers transact) than a

[63] Depending on the advertisers they serve, the ad buying tools can enable varying degrees of advertiser input into the bidding algorithm.
[64] *See* Google. "Determine a bid strategy based on your goals." Accessed on June 6, 2024. https://web.archive.org/web/20240602100502/https://support.google.com/google-ads/answer/2472725
[65] Some platforms also offer derivative services, such as shipping, fraud protection, *etc.*

bazaar (which largely serves as a meeting point for buyers and sellers to engage in whatever interaction they like) or Craigslist (which functions largely, although not entirely, like a digital bazaar).

78. An ad exchange's primary goal is to optimize its revenue, which is typically earned as a fraction of payments made when it facilitates a transaction between an advertiser and publisher.[66] Google operates an exchange called AdX, which was later combined with DFP into Google Ad Manager.[67]

79. In sum, publishers sell impressions, and use ad servers to manage this process. Advertisers buy impressions and use ad buying tools to manage this process. Exchanges intermediate this process by connecting publishers to advertisers. Figure 7 below presents the participants and the intermediaries in the online display ads market, with each Google intermediary.

Figure 7: Participants and intermediaries in the online display ads market

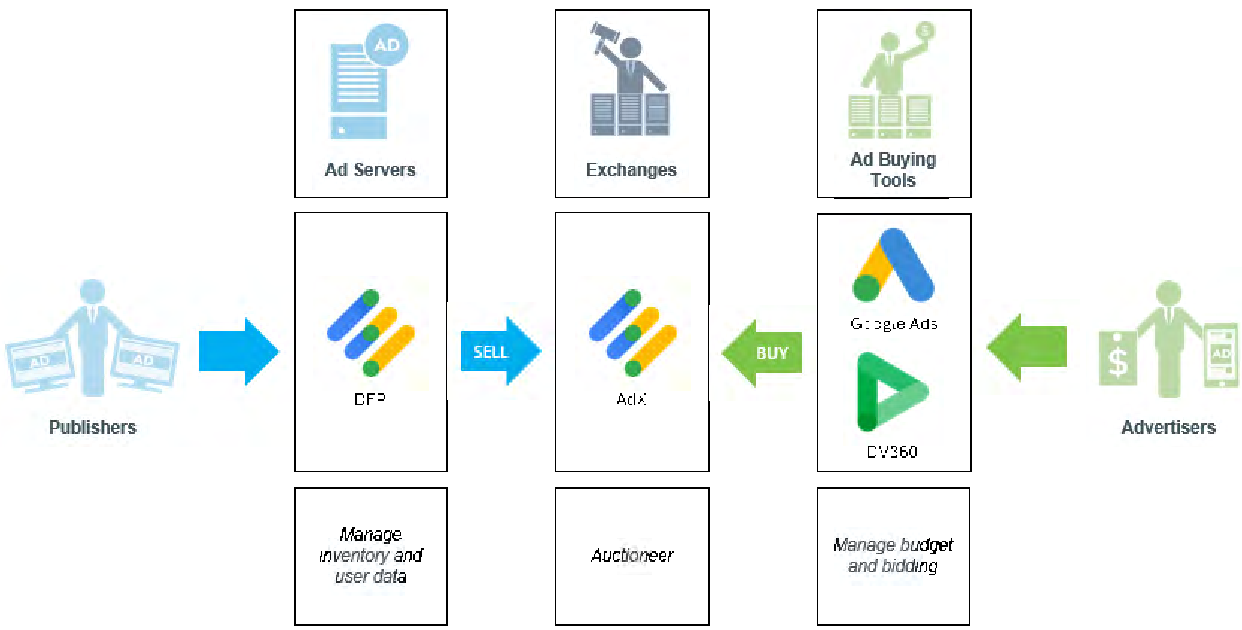

80. In order to better understand each key player's role in the markets for the display ads, consider the following example. The New York Times is a publisher that uses DFP as its ad server

[66] The conducts described in this report allege examples where AdX claims to earn revenue as a fraction of payments made when it facilitates a purchase, but actually collects revenue through a more complicated mechanism.
[67] *See* Jonathan Bellack. "Introducing Google Ad Manager" (June 27, 2018). Accessed on May 31, 2024. https://web.archive.org/web/20240303134019/https:/www.blog.google/products/admanager/introducing-google-ad-manager/

and has a direct deal with the advertiser Altra, a running shoe company. Altra pays the New York Times $10 per impression[68] for up to 1000 impressions per month to users over the age of 25 who enjoy running. When a user visits the New York Times, DFP is informed, and it decides how to sell the ad slot. DFP is aware of the direct deal with Altra and (if the user is over 25 and likes running) could display an Altra ad to help fulfill its direct deal. Alternatively, if for example the user did not fulfill the targeting criteria, DFP could contact various exchanges, such as AdX and PubMatic. AdX and PubMatic solicit bids from ad buying tools. AdX has access to the ad buying tool Google Ads, which is perhaps bidding for a local shoe seller, and DV360, which is bidding for Nike. PubMatic perhaps has access to DV360. At this point, both AdX and PubMatic are running their own auction for the impression. The bidders in this auction are the ad buying tools, who submit bids automatically on behalf of the advertisers.[69] AdX and PubMatic then execute their auctions, and pass on the resulting clearing prices to DFP, which then picks who gets to serve an ad. If that is AdX, and Nike is AdX's winning bidder, then Nike wins the impression. Along the way, DV360, AdX, and DFP each collect a fee for their services. For example, perhaps Nike wins AdX's auction at a price of $17. DV360 charges Nike $20 and takes 15% of Nike's $20 payment (a $3 fee) in order to pay $17 to AdX. AdX then takes 20% of DV360's $17 payment (a $3.40 fee) and pays $13.60 to the New York Times.[70] This example is illustrated in Figure 8 below.

[68] The numbers I provide in the examples throughout the report may not correspond to the frequently seen prices in real life. I instead choose numbers that enable the cleanest examples of the intended notions.
[69] Note that, depending on the particular auction format of the ad server, it may not be the case that both AdX and PubMatic are ever given the opportunity to submit a bid. This will be discussed further in this section under the waterfall.
[70] Importantly, DV360 and AdX collect a fee based on the particulars of this transaction (i.e., if the impression transacts at $10, they will collect a higher fee than if the impression transacts at $1). DFP collects a fee independent of the particulars of this transaction (i.e., perhaps the fee is a flat $.01 per transacted impression).

Figure 8: Nike gets the impression by winning the AdX auction through its ad buying tool DV360

B. Auction Formats Unique to Online Display Advertising

81. In Section II, I overviewed standard auction concepts such as the first- and second-price auctions, reserves, and personalized reserves. These concepts are all directly relevant to the case at hand. This subsection overviews two further concepts that are central to the case and prevalent in the online display advertising ecosystem: the waterfall and header bidding.

82. One building block for both is the concept of **line items**. Within an ad server, a line item refers to a potential demand source (e.g. an advertiser or an intermediary for an advertiser). Line items can be complicated and contain many different types of information, such as (a) an ad to display in case this line item is selected, (b) information about the buyer and where to look for payment in case this line item is selected, (c) the price that would be paid in case this line item is selected, (d) criteria that determine which impressions are permitted to select this line item or (e) other metadata to aid the process.[71]

83. The following are a few examples of how demand sources exist as line items.

a. A **guaranteed direct deal** from an advertiser offers price *p* per impression for any impression that satisfies proposed coarse targeting criteria. The advertiser expects exactly some number of impressions to be displayed per time period. When a

[71] *See* Google. "About line items." Accessed on May 31, 2024. https://web.archive.org/web/20240216155011/https:/support.google.com/admanager/answer/9405477?hl=en (Google Ad Manager's online documentation on line items.)

publisher creates a guaranteed direct deal as a line item, it includes the targeting criteria, the price per impression, and where to look for payment. The ad server tracks the numbers to ensure that the right number of impressions are sold.[72]

b. A **non-guaranteed direct deal** from an advertiser offers price *p* per impression for any impression that satisfies proposed coarse targeting criteria. The advertiser might also place a cap on the number of times this direct deal can be fulfilled per time period. When a publisher creates a non-guaranteed direct deal as a line item, that line item includes the targeting criteria, price per impression, and where to look for payment. The ad server tracks to ensure that the cap is not exceeded.[73]

c. An **individual exchange** is also a line item. When a publisher wants to elicit bids from a specific exchange, they add the exchange as a line item. Note that there is a distinct line item for each exchange.[74]

1) The Waterfall

84. The first key auction concept unique to the online display advertising ecosystem is the **waterfall**, which is a process used by an ad server to sell an impression. When selling an impression via the waterfall, the ad server visits line items one at a time in a priority order set by the ad server and the publisher.[75] When a line item is selected as the winner of the auction, the waterfall concludes.

85. Ad servers typically prioritize direct deals ahead of ad exchanges. They first check whether the impression meets the coarse targeting criteria for a direct deal, and if so allocate the

[72] Guaranteed direct deals would be a type of what is called a "sponsorship line item." *See* Google. "Sponsorship line items." Accessed on May 31, 2024.
https://web.archive.org/web/20221209041446/https:/support.google.com/admanager/answer/177426?hl=en
Some guaranteed direct deals can be based on a percentage of all the impressions satisfying the targeting criteria as well, such as "25% of impressions coming from women in Plano, TX."

[73] Non-guaranteed direct deals would be a type of what is called a "price priority line item" in Google Ad Manager's online documentation. *See* Google. "Price Priority line items." Accessed on May 31, 2024.
https://web.archive.org/web/20240216154933/https:/support.google.com/admanager/answer/79306?hl=en

[74] These are called "ad exchange line items" in Google Ad Manager documentation. *See* Google. "Ad Exchange line items." Accessed on May 31, 2024.
https://web.archive.org/web/20221012075051/https:/support.google.com/admanager/answer/188523?hl=en

[75] The ad server determines the general groups of line items and the ranking among these groups. *See* Google. "Line item types and priorities." Accessed on May 31, 2024.
https://web.archive.org/web/20240216154938/https:/support.google.com/admanager/answer/177279?hl=en
Within these groups. the publisher might manually set a line item order or rely on the ad server for automatic ordering based on some metric, such as the value CPM. *See* Google. "Value CPM." Accessed on May 31, 2024.
https://web.archive.org/web/20221202071803/https:/support.google.com/admanager/answer/177222?hl=en

impression to the highest paying direct deal without visiting any exchanges.[76] If the impression is not selected for any direct deal, the ad server then visits exchanges one at a time, in an order set by the publisher, with a price floor of r set by the publisher, possibly different for different exchanges.[77] When visited, the exchange can either claim the impression (and pay at least r to the publisher) or pass.[78] Figure 9 below illustrates such an example. The waterfall format was in use by DFP since before DFP was purchased by Google.[79] Since the rise of header bidding in the 2010s, which is discussed below, the waterfall has played a relatively smaller role than it once played.[80]

[76] Prioritization of direct deals over exchanges is a curious feature of the waterfall. If I were to design a waterfall-like format from scratch and I were unconstrained by technological challenges, I would (a) find the maximum payment v I could get from a direct deal for this impression (maybe $v = 0$, if it satisfies no direct deal targeting criteria), then (b) visit exchanges in the waterfall but setting reserves informed by v (for example, I would certainly never set a reserve lower than v, because I would rather just fulfill a direct deal for price v than sell to an exchange below v and I might certainly set a reserve above v, in order to get a chance at even greater revenue) and then (c) if the waterfall completed without any exchange paying, I would use the impression to fulfill the direct deal and collect my v for doing so. I would do this because if I decide to take a direct deal without visiting exchanges, I limit myself to exactly v, while if I instead decide to visit exchanges with all reserves higher than v before deciding on the direct deal, I guarantee myself at least v (because I can always fall back on the direct deal if all exchanges pass), but have a shot at more than v (if I get lucky and an exchange bids). Similarly, if I decide to pass on a direct deal without visiting exchanges, I may wind up with 0 if all exchanges pass, while if I instead visit exchanges first and they all pass, I can still get something via direct deal. The documentation I have access to does not explicitly state why the opposite decision was made, but surrounding context clues suggest this was likely due to technical limitations in a novel ecosystem using software initially developed for a simpler ecosystem. Anecdotal evidence also suggests that direct deals typically paid much more per-impression than live ad auctions anyway, and so therefore the loss due to suboptimal ordering may have been minimal. *See* Google. "Understand Direct and Programmatic Ad Revenue." Accessed on May 31, 2024.
https://web.archive.org/web/20231226200704/https:/newsinitiative.withgoogle.com/resources/trainings/grow-digital-ad-revenue/understand-direct-and-programmatic-ad-revenue/
There could certainly be other reasons too. Still, I do not mean to imply that this curiosity has a simple resolution. Because direct deals are made with coarser targeting criteria than live ad auctions, direct deal advertisers may insist on being considered first in order to avoid becoming victims of "cream-skimming" (where among impressions that satisfy the same coarse targeting criteria, live ad auctions win the "good" impressions and leave the "bad" ones for direct deals). I share this commentary primarily to give an example of how an auction theorist might reason through the process of auction design, although this particular thought process plays a role in my later analysis of Enhanced Dynamic Allocation.
[77] These reserves can be set by the publisher on their ad server, or on the ad exchange integrated into their ad servers. Also, there are third party tools who provide revenue optimization services, in some cases they may set the reserve prices on behalf of the publishers. Throughout the report, I mostly abstract away from these differences, since they are functionally the same.
[78] Payments do not occur instantaneously on a per-impression basis. Publishers get periodical payouts depending on the ad server they choose to work with. *See* Google. "Ad Manager payment timelines." Accessed on May 31, 2024.
https://web.archive.org/web/20221209003032/https:/support.google.com/admanager/answer/2671030?hl=en (Google's ad server payout options.)
[79] GOOG-NE-10780865 at -78, 79. May 5, 2020. "Clearing Up Misconceptions About Google's Ad Tech Business."
[80] *See* Sarah Sluis. "The Rise Of 'Header Bidding' And The End Of The Publisher Waterfall" (June 18, 2015). Accessed on May 31, 2024.
https://web.archive.org/web/20240216154913/https:/www.adexchanger.com/publishers/the-rise-of-header-bidding-and-the-end-of-the-publisher-waterfall/

Figure 9: An impression is allocated to AdX because it does not satisfy the targeting criteria for a high paying Altra direct deal

86. To understand how the waterfall process works and how it relates to common auction formats, imagine an impression arrives for a user over the age of 25 who likes running. DFP notices that this satisfies the targeting criteria for a direct deal with Altra, and decides to display Altra's ad. The waterfall terminates here.[81] This example is illustrated in Figure 10 below.

Figure 10: An impression is allocated to a direct deal since it satisfies the targeting criteria

[81] In other words, the selling process stops because the impression is sold.

87. Another new impression arrives for a user over the age of 25 who likes expensive handbags. DFP does not have a direct deal that meets this targeting criteria, and so moves on to visit exchanges. It first visits PubMatic, for which the publisher set a reserve of $5. PubMatic then solicits bids through ad buying tools and does not find a sufficiently high bid and passes. Next, DFP visits AdX, for which the publisher sets a reserve of $4. AdX runs a second-price auction, soliciting bids through demand side platforms and finds that the ad buying tool for Gucci submitted a bid of $8, and the ad buying tool for Louis Vuitton submitted a bid of $6. This clears AdX's auction at $6, AdX takes a 20% cut, which puts the maximum reserve they can beat at $4.80. AdX claims the impression. Gucci gets to display their ad, paying $6 to AdX who charges $1.20 and pays $4.80 to the ad server.[82] This example is illustrated in Figure 11 below.

Figure 11: Gucci wins an impression through AdX because the impression does not satisfy the targeting criteria for a high paying Altra direct deal

88. These examples emphasize a few aspects of waterfall auctions: (a) the decision of whether to use a direct deal or not occurs before the ad server visits exchanges, (b) an exchange that is late in the waterfall may never have an opportunity to bid at all, if the impression is sold earlier in the waterfall, and (c) advertisers can set different reserves for different exchanges.

[82] This example abstracts away from the fee that the ad buying tool would charge on top of the clearing price, for the sake of clarity.

89. DFP gives publishers the option to set whatever reserves they like.[83] Documentation suggests that most publishers chose to set reserves based on the historical average payment from that exchange.[84] Note that, given access to the distribution of historical payments, it is suboptimal to simply set the average payment as the reserve. This holds even when there is just a single exchange. Seminal work of Myerson (1981)[85] describes the revenue-maximizing reserve in this case. With multiple exchanges visited through the waterfall, optimal reserves are even more complex.

90. Documentation also suggests that publishers sort exchanges primarily in decreasing order of historical average payment, although exchange fill rate (the fraction of offers an exchange accepts) and ad quality (the quality of the visual displayed in the auctioned ad space) might play a role too.[86] This heuristic of setting the waterfall ordering on the basis of historical CPM[87] averages generates nontrivial incentives for the exchanges. A higher winning bid by exchanges increases their future reserves (because their average bid goes up), making future impressions more expensive. On the other hand, a higher reserve puts the exchange earlier in the waterfall, which gives them access to more impressions. Sorting exchanges in decreasing order of reserve is natural since (a) combined with the simple average CPM heuristic for setting reserves, exchanges have the indirect opportunity to pay more to be placed earlier in the waterfall, and (b) if instead exchanges eschew this opportunity and simply pay exactly the reserve to prevent the

[83] They were later constrained in their freedom to choose appropriate reserve prices by Unified Pricing Rules conduct, which I discuss in Section VI.
[84] GOOG-NE-10780865 at -81. May 5, 2020. “Clearing Up Misconceptions About Google’s Ad Tech Business.” (“Publishers typically set the net value CPM for their booked static remnant line items based on their estimates of what CPM the line item would likely generate (taking into account its historical performance) or based on a fixed-price the publisher had negotiated with a particular remnant demand partner.”)
[85] Roger B. Myerson. “Optimal Auction Design.” *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73.
[86] *See* GOOG-NE-10780865 at -79. May 5, 2020. “Clearing Up Misconceptions About Google’s Ad Tech Business.” (¶3 describes how publishers ordered demand sources.)
[87] CPM refers to “cost per mille.” It is the cost of purchasing one thousand impressions for an advertiser.

reserve price from increasing in future, the revenue-optimal ordering for the publishers is indeed to sort exchanges in decreasing order of reserve.[88, 89]

91. From an auction theory perspective, waterfalling as a procedure for selling impressions is inefficient and sacrifices revenue. This is because waterfalling forces the ad server to decide whether to sell to one exchange before learning what other exchanges might bid. A first- or second-price auction among exchanges (as opposed to the waterfalling) simultaneously considers all exchanges, which avoids this problem. Indeed, there are several research papers quantifying the suboptimality of sequential auction formats (specifically, posted-price mechanisms,[90] which are the closest common auction format to waterfalling) as compared to simultaneous auction formats (specifically, second-price auctions).[91] Moreover, the concept of sequential decision-making with partial information versus simultaneous decision-making with all information is well-studied within computer science broadly under the field of Online Algorithms (here, “online” refers to “making binding decisions one at a time with incomplete information” rather than “on the internet”),[92] and it is well-understood that binding sequential decisions come at a loss compared to a single decision with all the available information.

92. An ad server might still employ the waterfall process even with these suboptimalities. One reason computer scientists study online algorithms is because making a single decision with full

[88] This leads to the question of why exchanges ever pay more than the reserve price. First, some exchanges may have inflexible contracts with advertisers that preclude them from being particularly strategic with bids. For example, perhaps an exchange agrees to take exactly a 20% cut of the winning bid and give the rest to the publisher. Then, if this exchange has an advertiser willing to pay up to $6, and their reserve is $4, the highest revenue the exchange can collect is by paying $4.80 to the ad server and collecting $1.20 on $6. If instead the exchange were to pay exactly $4, this would correspond to taking a 20% cut of $5, which is just $1. However, this reasoning does not apply if exchanges were not bound by such contracts. An unbound exchange in this example could pay $4 to the ad server, collect $6 from the advertiser, and pocket $2. Second, if the ad server indeed uses the average of past clearing prices as a heuristic to set reserves, exchanges’ incentives are complicated. If other exchanges are likely to purchase the impression, being early in the waterfall is the only way to get a shot at the impression. Therefore, exchanges may wish to submit higher bids than what is needed to win in order to increase their average historical bid and move earlier in the waterfall.

[89] The waterfall can be analyzed in comparison to the other common auction formats. The closest standard auction format to waterfalling is a posted-price mechanism. In a posted-price mechanism, exchanges would be visited one at a time, offered the impression at a personalized reserve, and could either pay the reserve to receive the impression, or pass. In a waterfall, the only difference is that exchanges can choose to pay above the reserve.

[90] In general, under posted-price mechanisms, the first bidder that clears their reserve is awarded the auctioned item.

[91] *See, e.g.*, Hartline and Roughgarden. “Simple versus Optimal Mechanisms.” *Proceedings of the 10th ACM Conference on Electronic Commerce*. 2009. pg. 225-234; Alaei, Hartline, Niazadeh, Pountourakis, and Yuan. “Optimal auctions vs. anonymous pricing.” *Games and Economic Behavior* vol. 118. 2019. pg. 494-510; Jin, Lu, Qi, Tang, and Xiao. “Tight Approximation Ratio of Anonymous Pricing.” *Proceedings of the 51st Annual ACM SIGACT Symposium on the Theory of Computing*. 2019. pg. 674-685.

[92] *See, e.g.*, Sleator and Tarjan. “Amortized efficiency of list update and paging rules.” *Communications of the ACM* vol. 28, no. 2. 1985. pg. 202-208; Krengel and Sucheston. “On semiamarts, amarts, and processes with finite value.” *Probability on Banach spaces*. 1978. pg. 197-266; Martin L. Weitzman. “Optimal Search for the Best Alternative.” *Econometrica* vol. 47, no. 3. 1979. pg. 641–654.

information is impossible, since some information is going to be revealed in the future.[93] This justification does not apply in this case since there is no unknown future information. Another reason why posted-price mechanisms are sometimes used in practice is because their simplicity makes it especially easy for users to interact, due to its take-it-or-leave-it nature. Some sellers decide that a pleasant buyer experience with exceptionally simple auctions outweighs the additional revenue they would earn from a first- or second-price auction.[94] This justification likely does not apply in this case either since the entities that bid in publishers' auctions are sophisticated exchanges whose entire business model is centered around optimally conducting auctions.[95] A third possible reason is technological challenges. Because the online display advertising ecosystem is complex and evolved in a piecemeal fashion, ad servers may have used infrastructure that was initially developed for a simpler ecosystem without exchanges.

93. Still, there are idiosyncratic reasons why the waterfall may be favored over other auction formats, such as technological limitations and/or simplicity of interactions elsewhere in the ecosystem.[96] However, this also suggests that if these idiosyncratic reasons could be addressed (e.g. with technological developments and/or comfort with complex interactions elsewhere in the ecosystem), it would instead make sense to use a simultaneous form that generates increased revenue.

2) Header Bidding

94. **Header bidding** refers to a simultaneous auction technology that was developed in the early 2010s.[97] Exchanges with real time bidding capability were present before,[98] but header bidding enabled these exchanges to compete in an auction of auctions for the impression. For header bidding, publishers embed a piece of code in the "headers" of their webpages and this code executes a first-price auction with personalized reserves[99] by soliciting bids from integrated

[93] For example, in the so-called "ski rental problem," a sequential solution decides before each ski trip whether to rent or buy skis. Here, a "full information solution" would need to know exactly how many times you intend to ski over your lifetime in order to determine whether renting or buying is more efficient. This solution is not an option because it is not possible for a person to know that information.
[94] *See, e.g.*, Shengwu Li. "Obviously Strategy-Proof Mechanisms." *American Economic Review* vol. 107, no. 11. 2017. pg. 3257-3287. (subsequent research is also relevant.)
[95] Of course, just because the exchanges themselves are extremely sophisticated does not necessarily mean that all publishers and all advertisers are comparably sophisticated.
[96] For example, in waterfalling, if an exchange is ever visited, they can promise the buyers that they will get that impression if they win with certainty. In simultaneous formats, they cannot.
[97] GOOG-NE-10780865 at -82, 83. May 5, 2020. "Clearing Up Misconceptions About Google's Ad Tech Business."
[98] *See* Interactive Advertising Bureau. "OpenRTB." Accessed on June 4, 2024. https://web.archive.org/web/20240326073202/https://www.iab.com/guidelines/openrtb/
[99] More specifically, a "floor price module" in Prebid enables personalized floors. *See* Oscar Bole. "How to make sure your floor setup is flawless?" (June 27, 2022). Accessed on May 31, 2024.

exchanges[100] before the webpage calls the ad server for the available ad spaces. If an exchange wins the auction at price *p*, this is passed on to the ad server as a line item. That is, the winning exchange does not immediately win the impression at price *p*, but the ad server learns that it is an option to sell this impression at price *p*. Header bidding is a free and open-source project developed by industry actors.[101] Figure 12 below illustrates the selling of an impression where the publisher makes the impression available to header bidding as well.

Figure 12: The selling of an impression from a publisher who utilizes a header bidding setup

95. To better understand how header bidding works, imagine that an impression arrives for a user over the age of 25 who likes running. Before the ad server executes, header bidding solicits bids from exchanges. The highest bidder that exceeds their personalized reserve is Nike through OpenX[102] for $4 and this is entered as a line item in the ad server. The ad server then begins the waterfall, notices that this impression satisfies the targeting criteria for a direct deal with Altra and decides to display Altra's ad. The waterfall terminates here (that is, although the header bidding

https://web.archive.org/web/20231202022909/https:/www.pubstack.io/topics/implementing-floor-prices#1; GOOG-TEX-00843142 at -45. September 3, 2019. "First-price bidding Update." ("the header bidders are called first and a first-price auction is run amongst them.")
[100] Sometimes, ad buying tools or ad networks directly bid in the header bidding auctions as well.
[101] GOOG-NE-10780865 at -83. May 5, 2020. "Clearing Up Misconceptions About Google's Ad Tech Business."
[102] OpenX is an ad exchange.

auction takes place, the resulting line item is not always visited by the ad server). This example is illustrated in Figure 13 below.

Figure 13: An impression is allocated to an Altra direct deal since it is higher priority than non-guaranteed line items

96. Now imagine another impression arrives for a user over the age of 25 who likes expensive handbags. Before the ad server executes, the header bidding solicits bids from exchanges. The highest bidder that exceeds their personalized reserve is Prada for $4.50 through OpenX and this is entered as a line item in the ad server. The ad server then begins the waterfall. The ad server does not have a direct deal that meets this targeting criteria, so it moves on to other line items. It first visits PubMatic with a reserve of $5. PubMatic then solicits bids through ad buying tools, does not find a sufficiently high bid and passes. The ad server next visits the header bidding generated line item for OpenX at $4.50 and sells the impression to OpenX winner Prada. Other exchanges ranked below, such as AdX, are not visited. This example is illustrated in Figure 14 below.

Figure 14: An impression is allocated to Prada who won the header bidding auction through OpenX

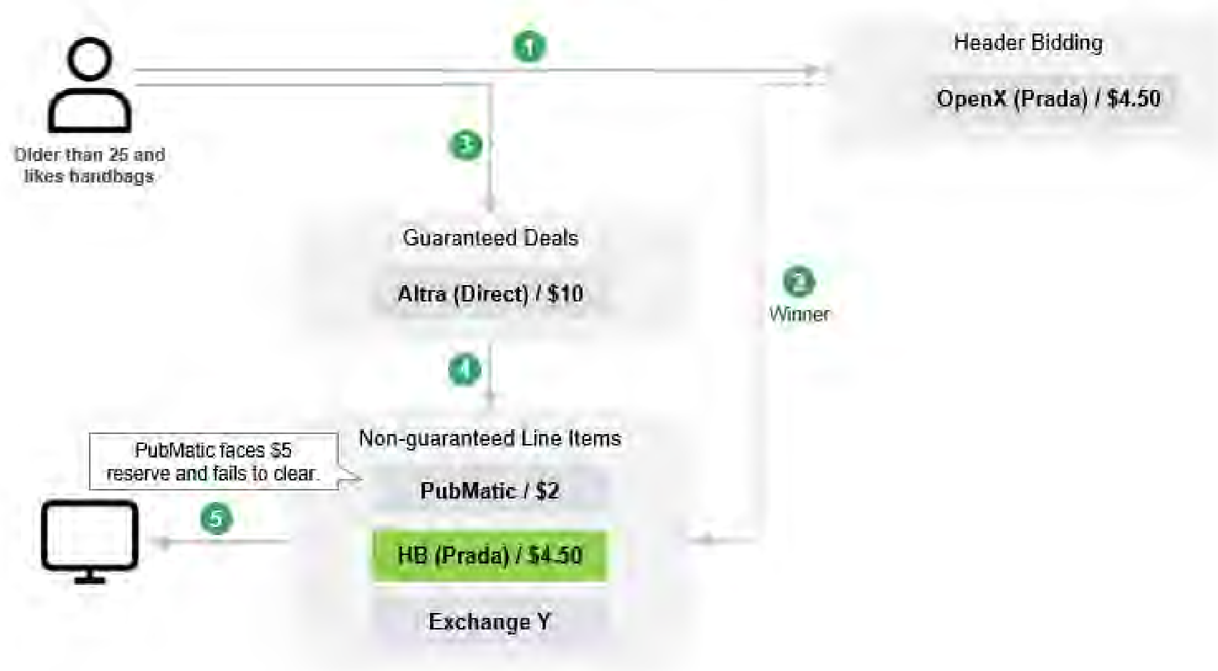

97. Lastly, imagine a new impression arrives for a user over the age of 25 who likes sports. Before the ad server executes, header bidding solicits bids from exchanges. The highest bidder that exceeds their personalized reserve is Dick's for $6 through OpenX and this is entered as a line item in the ad server. The ad server then begins the waterfall process. The ad server does not have a direct deal that meets this targeting criteria, and so it moves on to other line items. It first visits the header bidding generated line item for OpenX at $6 and sells the impression to Dick's. No other exchanges ranked below, such as PubMatic or AdX, are visited. This example is illustrated in Figure 15 below.

Figure 15: An impression is allocated to Prada who won the header bidding auction through OpenX

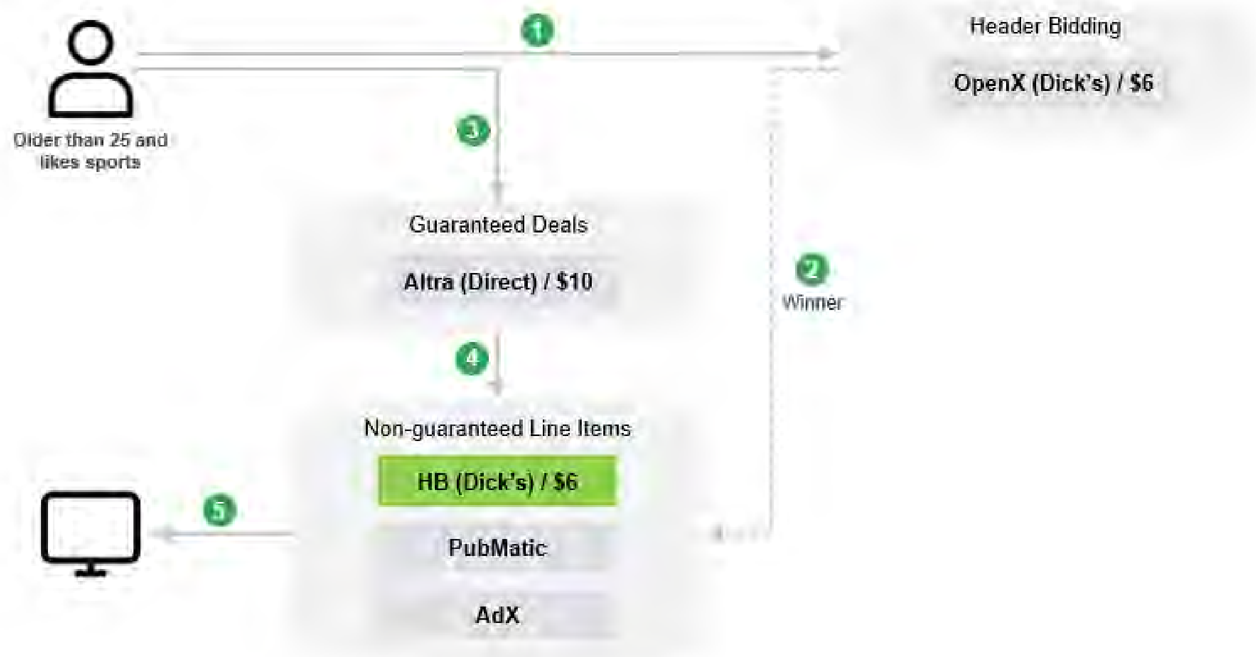

98. The examples above emphasize a few aspects of the logistical role of header bidding in the waterfall such as (a) the header bidding auction executes before the ad server, and therefore is always executed (even if the header bidding winner does not ultimately win the impression), (b) the winning header bidder generates a line item in the waterfall, (c) the header bidding line item is typically lower on the waterfall ranking than the direct deals are and (d) the publisher can prioritize the header bidding line item relative to each exchange based on the winning bid.

a) Comparison of Header Bidding and the Waterfall

99. Header bidding addresses the waterfall inefficiencies by soliciting bids from all exchanges simultaneously and running a first-price auction with personalized reserves, rather than visiting exchanges sequentially, which risks both selling the impression too early (and missing out on a better price later) and selling the impression too late (foregoing a good deal early in hopes of finding a better one later that never materializes). Judging by header bidding's rapid adoption rate among publishers,[103] it seems that header bidding successfully addressed these inefficiencies

[103] GOOG-TEX-00105361 at -85. April 28, 2017. "FAN Bidding in to DPI and AdMob." (provides metrics for header bidding adoption among publishers, and the document points out the publisher adoption of header bidding "shows risk is increasing" for Google.)

and increased publisher revenue, which is expected from an auction theory perspective, as explained above.

100. The amount of data collected is another benefit of a simultaneous auction format like header bidding over a sequential format like waterfalling. Specifically, in a simultaneous auction format, the header bidding setup records the bids of every exchange whereas in a sequential auction format the ad server learns bids of only the winning exchange and those before it in the waterfall. As noted in previous sections, data is crucial in auctions in order to set optimal reserves.[104, 105]

IV. CONDUCT ANALYSIS: DYNAMIC ALLOCATION AND ENHANCED DYNAMIC ALLOCATION

101. In this section, I analyze the conduct referred to as Dynamic Allocation and later updated to Enhanced Dynamic Allocation.[106] In this section, I draw conclusions regarding their effects. I find that Dynamic Allocation led to a higher win rate[107] and higher revenue for AdX as well as a lower win rate and lower revenue for non-Google exchanges. Furthermore, Enhanced Dynamic Allocation led to an increase in win rate and increase in revenue for AdX and reduced the value of direct deals for advertisers. Reducing the value of direct deals for advertisers should decrease the revenue earned by publishers via direct deals.

102. Moreover, any related conduct that causes AdX to clear publisher-set reserves more often exacerbates many of my conclusions. Dynamic Revenue Sharing is one such conduct, which I describe in Section VII.[108] Throughout, I briefly note conclusions whose magnitude is increased due to Dynamic Revenue Sharing and related conducts.

---

[104] This is especially true in ad auctions, due to the fact that each impression is unique, and the task of predicting an advertiser's/exchange's bid for a novel impression must be inferred by their past bids on similar but not necessarily identical impressions.

[105] This particular form of additional data is commonly studied in the sub-field of machine learning called "online learning." Learning every exchange's bid is called "expert feedback" and learning only the winning exchange's bid is called "bandit feedback." It is well-understood, including quantitatively, that expert feedback enables faster and more accurate learning than bandit feedback. *See, e.g.,* Aleksandrs Slivkins. "Introduction to Multi-Armed Bandits" (November 2019). https://arxiv.org/pdf/1904.07272 (chapters 5 and 6.)
Header bidding produces expert feedback, while waterfalling produces feedback somewhere between bandit and expert, depending on how late the winning exchange is on the waterfall.

[106] There are many descriptions of Dynamic Allocation and Enhanced Dynamic Allocation, according to the documents I have reviewed. The descriptions I provide here reflect my best understanding. It is possible that at some points in time Dynamic Allocation and Enhanced Dynamic Allocation worked slightly different compared to the descriptions I provide here.

[107] I use the term "win rate" to refer to the number of impressions won divided by the number of impressions made available in the open web display ads market.

[108] Briefly, Dynamic Revenue Sharing is a program where AdX sometimes lowered its take rate in order to clear impressions for which it did not solicit a large enough bid to clear the publisher's price floor plus its standard take rate. See Section VII for detailed description.

A. Dynamic Allocation

103. The impact of Dynamic Allocation on the waterfall process depends on the types of line items present in the waterfall. More specifically, whether there are only static demand sources or there are both static and live demand sources[109] affects how Dynamic Allocation works, as well as its impact on the auction procedure and outcomes. Hence, I analyze Dynamic Allocation separately for static demand sources and live demand sources.

1) Dynamic Allocation with static demand sources

104. I first present an overview of Dynamic Allocation during the period when it was first introduced.[110] Initially, all line items were static, so Dynamic Allocation addressed a natural shortcoming of the waterfall format. When all line items competing with AdX are static, **Dynamic Allocation with Static Line Items** adjusts the waterfall process in the following manner:

a. First, Google's ad server DFP processes the high priority line items[111] that are not affected by Dynamic Allocation (such as direct deals). If any high priority line item succeeds, the impression is sold, and the waterfall terminates without continuing to subsequent steps.

b. Every low priority line item, including AdX, has both a price floor and a Value CPM.[112] Next, DFP selects the highest Value CPM among all low priority static line

[109] Throughout this report, I use the term "static line item" when referring to line items that do not correspond to outcomes of any auctions. For example, the line items that Google documentation refers to as sponsorship or standard would be static line items. *See* Google. "Line item types and priorities." Accessed on May 31, 2024. https://web.archive.org/web/20240216154938/https:/support.google.com/admanager/answer/177279?hl=en
I use the term "live demand sources" when referring to the ad exchange line items. They are "live" since they hold an auction before submitting a clearing price.

[110] Dynamic Allocation was introduced by DoubleClick, prior to Google's purchase of the company. DoubleClick documentation from that time points to 2007 as the introduction of Dynamic Allocation. *See* Google. "DoubleClick Advertising Exchange." Accessed on May 31, 2024.
https://web.archive.org/web/20071001100309/http:/www.doubleclick.com/products/advertisingexchange/index.aspx
Google documentation claims it is 2008, while agreeing that it predates Google's acquisition of DoubleClick. GOOG-AT-MDL-008991406 at -6. ("2008—Dynamic Allocation […] pre doubleclick acquisition.")

[111] Throughout the report, I use "guaranteed" and "high priority" interchangeably when referring to line items that are at priority 1-10 by Google's standards. Similarly, I use "non-guaranteed" and "low priority" interchangeably when referring to line items that are at priority 12-16. *See* Google. "Line item types and priorities." Accessed on May 31, 2024. https://web.archive.org/web/20240216154938/https:/support.google.com/admanager/answer/177279?hl=en
This is in line with the terminology that Google uses. A Google engineer stated that "[redacted] Declaration of Nitish Korula."

[112] Value CPMs are set by the publishers, and they usually correspond to the value of those line items for the publishers. Google provides the following formula to estimate the value CPM: Value CPM = (Total revenue received from ad tags associated with selected line item/Total number of impressions Ad Manager sent to the selected line

items, that satisfy the targeting criteria for this impression and stores this value as a reserve price *r*. The publisher sets each Value CPM according to Google's documentation, and equal to the expected value derived from each line item. Because all line items are static, the reserve price then captures the expected maximum value that the publisher can earn by selling the impression to a low priority line item.[113, 114]

c. Next, DFP calls AdX to run an auction with a reserve price equal to the maximum of *r* and AdX's price floor. If AdX returns a price higher than this reserve, the impression is sold through AdX, and the waterfall terminates without continuing to subsequent steps.

d. Finally, if AdX fails to return a bid higher than this reserve, the impression is sold to the low priority demand source with the highest Value CPM (that satisfy the targeting criteria for this impression).[115]

---

item)*1000. Google. "Value CPM." Accessed on May 31, 2024. https://web.archive.org/web/20221202071803/https:/support.google.com/admanager/answer/177222?hl=en A Google engineer stated "When configuring a remnant line item, a publisher must specify a rate (i.e., a price) for the line item. The publisher may also specify a Value CPM; this might be done when the price does not accurately reflect the value to the publisher of serving the line item. For example, if the publisher gave an advertiser a discount on a line item, the publisher could enter the discounted price as the rate and the undiscounted price as the Value CPM. If a publisher does not specify a Value CPM, then Google sets the Value CPM equal to the rate." and "By configuring multiple line items with different targeting criteria, publishers could configure different Value CPMs based on, for example, time of day or the geography of the relevant user, even for the same demand partner. Some publishers set Value CPMs based on their estimates of what CPM a line item would likely generate (taking into account its historical performance) or based on a fixed price the publisher had negotiated with a particular remnant demand partner. Some publishers set Value CPMs higher than their estimates of what CPM a line item would likely generate to increase competitive pressure in the AdX auction or for other reasons. Under Dynamic Allocation, the Value CPM associated with the best eligible non-guaranteed line item could set the reserve price in the AdX auction." GOOG-AT-MDL-008842393 at -96. August 4, 2023. "Declaration of Nitish Korula."

[113] These static line items can pay more or less than the assigned Value CPM.

[114] A sophisticated publisher could ignore Google's suggested formulas and set the Value CPMs however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default. Throughout the text, I use the term "sophisticated publisher" with the following context in mind: Publishers can have varying levels of sophistication when attempting to optimize their revenue in the online display advertising ecosystem. On one end, a 'typical' publisher may set parameters according to their ad server's suggested text without developing a detailed understanding of how those parameters are used. At the other end, a 'sophisticated' publisher may fully digest all available documentation and aim to optimize parameters based on their use case, ignoring suggested text. They may even be able to optimize while accounting for the possibility of conduct that is never disclosed in publicly available documentation. Furthermore, I use the term "sophisticated advertiser" with the following context in mind: Advertisers have varying levels of sophistication when attempting to optimize their outcomes in the online display advertising ecosystem. On one end, a 'typical' advertiser may trust their ad buying tool to optimize on their behalf and input correct information whenever requested (i.e., a 'typical' advertiser would simply input their correct value for an impression when asked). At the other end, a 'sophisticated' advertiser may fully digest all available documentation and aim to optimize inputs to their ad buying tool based on how these inputs are used, ignoring the ad buying tool's recommendations. They may even be able to optimize while accounting for the possibility of conduct that is never disclosed in publicly available documentation.

[115] An overwhelming majority of Google's documentation, and their testimony in GOOG-NE-10780865, supports the given definition. GOOG-NE-10780865 at -81. May 5, 2020. "Clearing Up Misconceptions About Google's Ad Tech

105. The excerpt in Figure 16 from a public Google report from 2010[116] reiterates how Dynamic Allocation with Static Line Items work. The document notes that "[Dynamic Allocation] uses this CPM value as the minimum CPM for the auction,"[117] as I described above.

Business." ("If AdX buyers did not bid above the floor price (on a net basis, i.e., after consideration of AdX's revenue share), the static remnant line item with the highest fixed or estimated price would win the impression.") However, some internal documentation supports the following conflicting definition: [redacted] Of course, any precise mathematical analysis differs between these two formats, although my relevant conclusions do not qualitatively differ. I focus my analysis on the definition which is overwhelmingly supported by Google's documentation.

[116] Google. "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers." Accessed on May 31, 2024. https://web.archive.org/web/20120130063019/http:/static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en/us/doubleclick/pdfs/DC_Ad_Exchange_WP_100713.pdf (pg. 3 describes mechanics of AdX under Dynamic Allocation.)

[117] Google. "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers." Accessed on May 31, 2024. https://web.archive.org/web/20120130063019/http:/static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en/us/doubleclick/pdfs/DC_Ad_Exchange_WP_100713.pdf

Figure 16: Excerpt from a public Google report explaining how Dynamic Allocation with Static Line Items works[118]

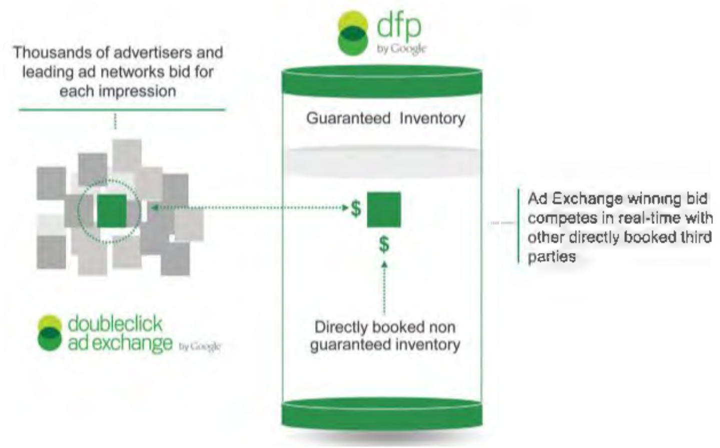

Dynamic allocation passes to the Ad Exchange the CPM value associated with the ad that the primary ad server has selected and is about to serve. The technology then uses this CPM value as the minimum CPM for the auction. If the Ad Exchange can provide the publisher with a net CPM value higher than they would have gotten from delivering their directly booked, non-guaranteed ad, the Ad Exchange will deliver an ad. If, however, the directly booked ad's CPM value is higher, it ignores any bids coming in from the Ad Exchange. As a result of this ad server integration, publishers essentially have a risk-free way to get the highest yield for every non-guaranteed impression they sell through their direct and indirect sales channels. An additional benefit of Dynamic Allocation is that it ensures there is a deep pool of ads to deliver for any given piece of inventory, reducing the probability that the publisher delivers a house or zero-value ad.

106. The following excerpt in Figure 17 from an internal Google slide deck[119] is also consistent with my explanation on how Dynamic Allocation with static line items works.

[118] Google. "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers." Accessed on May 31, 2024. https://web.archive.org/web/20120130063019/http:/static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en/us/doubleclick/pdfs/DC_Ad_Exchange_WP_100713.pdf (pg. 3 describes mechanics of AdX under Dynamic Allocation.)

[119] GOOG-NE-08112779. "PBS Basics Training (3) AdX Basics."

Figure 17: An excerpt from an internal Google slice deck explaining Dynamic Allocation with Static Line Items compares the highest valued line item to the AdX clearing price[120]

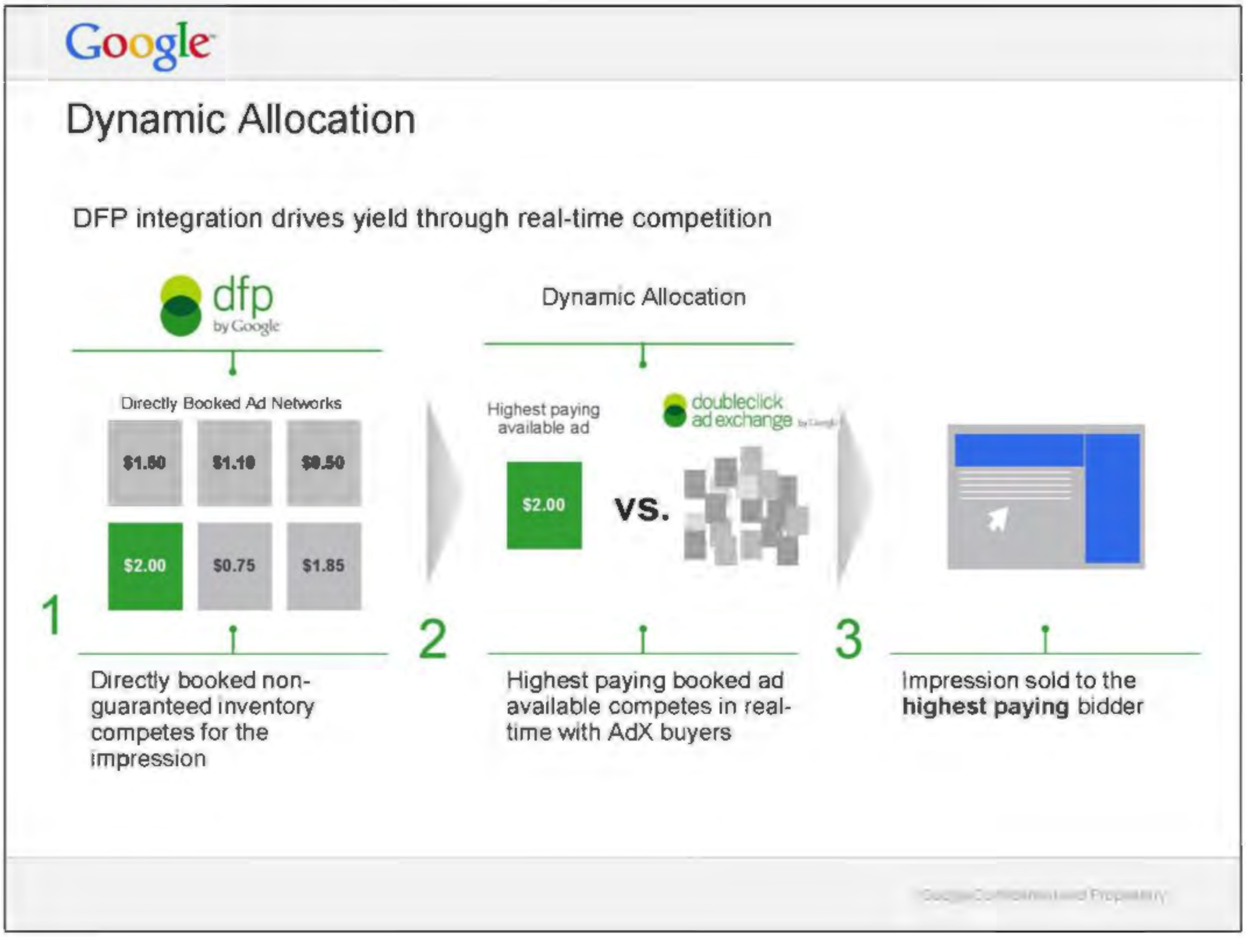

107. Several internal Google documents provide details on how Dynamic Allocation worked. An internal Google document states that the highest available bid information is passed on to AdX as a "dynamic floor price."[121] A publicly available Google document states that "Dynamic Allocation is a unique technology that works by passing to the Ad Exchange the CPM value associated with any non-guaranteed ad that DFP is about to serve…the Ad Exchange only serves ads when it can offer a higher price for ad space."[122]

[120] GOOG-NE-08112779 at -94. "PBS Basics Training (3) AdX Basics."
[121] GOOG-NE-03597611 at -19. December 15, 2011. "Mysteries of Dynamic Allocation."
[122] Google. "Maximizing advertising revenues for online publishers." Accessed on May 31, 2024. https://web.archive.org/web/20160911040651/https:/static.googleusercontent.com/media/www.google.com/en/googleblogs/pdfs/revenue_maximization_090210.pdf

108. To illustrate how Dynamic Allocation with Static Line Items works, imagine an impression arrives for a user over the age of 25 who likes running. DFP notices that this satisfies the coarse targeting criteria for a direct deal with Altra, displays Altra's ad, and the waterfall (even with Dynamic Allocation) ends here. This example is illustrated in Figure 18 below.

Figure 18: An impression is allocated to an Altra direct deal in Dynamic Allocation because it fulfills the targeting criteria

109. Another new impression arrives for a user over the age of 25 who likes expensive handbags. DFP does not have a direct deal that meets this targeting criteria, and so moves on to lower-priority line items and observes that the highest Value CPM option is from a static line item for $4. DFP then calls AdX with a reserve of $4. AdX does not find a buyer above $4, and the waterfall concludes by allocating the impression to the static line item for $4. This example is illustrated in Figure 19 below.

Figure 19: An impression is allocated to a non-guaranteed line item because the impression fails to clear the targeting criteria of an Altra direct deal, and AdX fails to clear the Dynamic Allocation reserve[123]

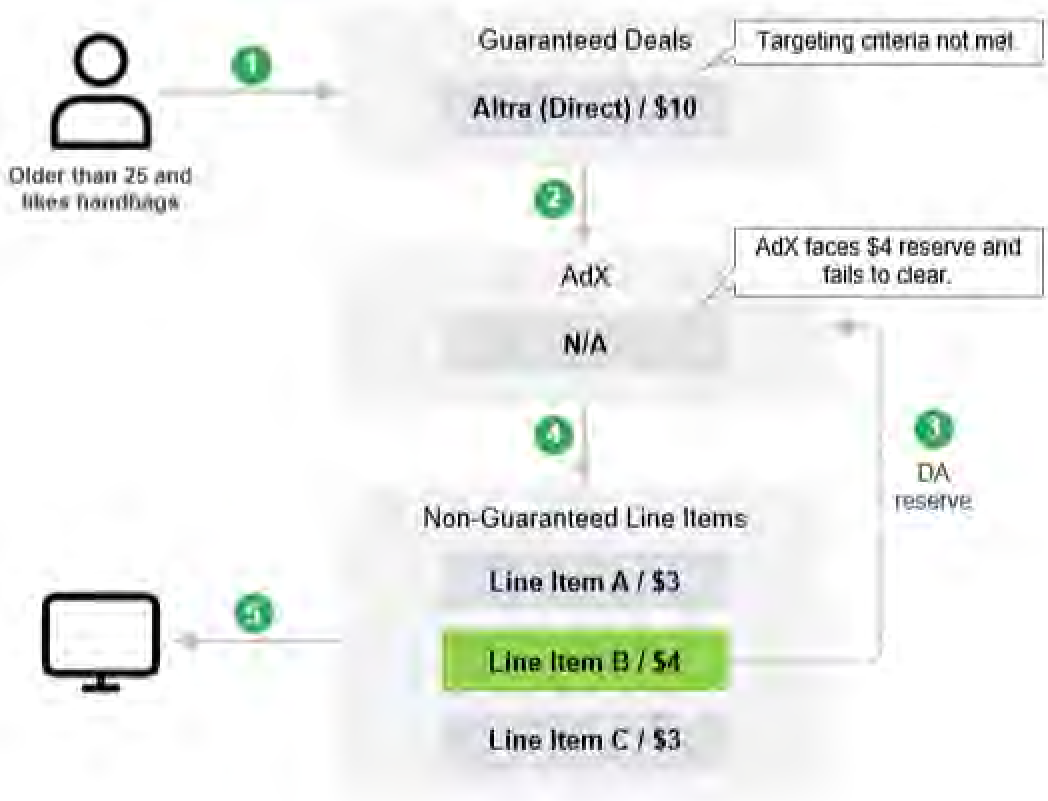

110. Lastly, an impression arrives for a user over the age of 25 who likes expensive handbags, and with further fine-grained cookies noting that the user has recently visited several handbag-shopping websites and lives in a wealthy neighborhood. DFP does not have a direct deal that meets the coarse targeting criteria of users over the age of 25 who like expensive handbags, and so it moves on to lower priority line items and observes that the highest Value CPM option is from a static line item for $4. DFP then calls AdX with a reserve of $4. AdX's auction concludes with Gucci winning at a clearing price of $6. Gucci wins the impression through AdX, paying $6.[124] This example is illustrated in Figure 20 below.

[123] In the figures, Dynamic Allocation is referred to as "DA."
[124] This example abstracts away from the role of ad buying tools, ad exchange fees and a publisher-set reserve price for AdX for clarity.

Figure 20: Gucci wins an impression through AdX because it is able to clear the Dynamic Allocation reserve

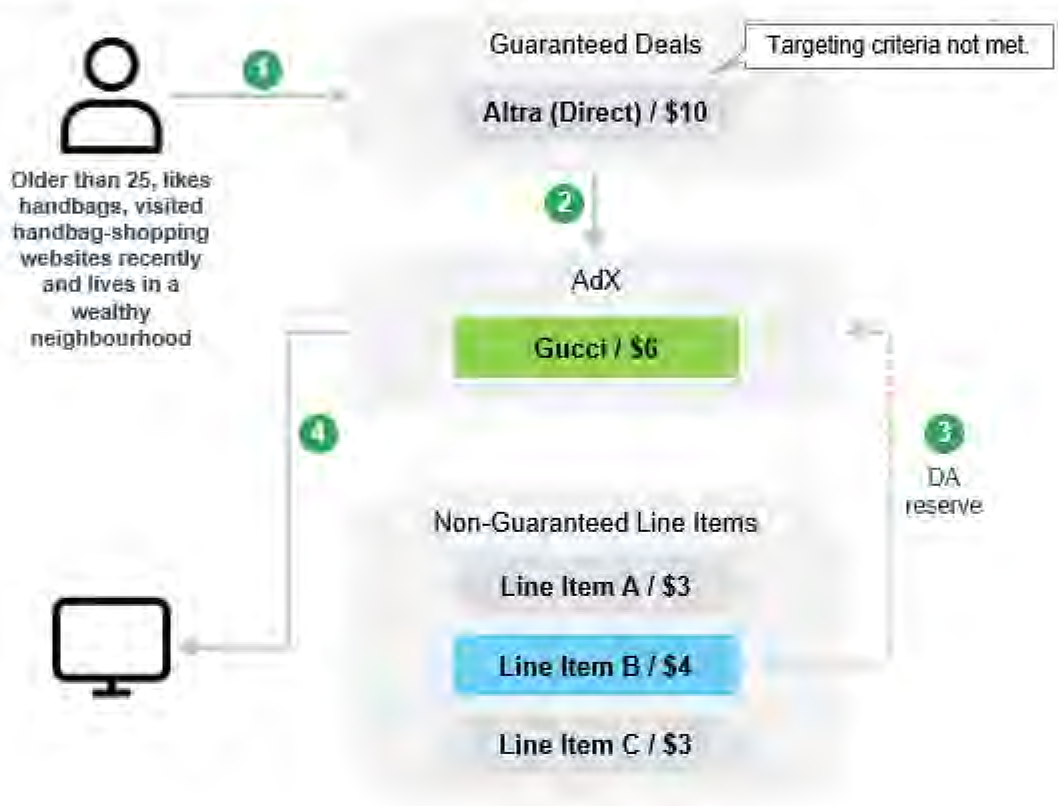

111. The waterfall format is a suboptimal mechanism for soliciting bids from live sources.[125, 126] The optimal format would solicit bids from all sources simultaneously and conclude the auction according to Myerson's (1981)[127] optimal auction.[128] In the special case where only AdX (or another single exchange) provides live demand and all other demand sources are static, Myerson's optimal auction is equivalent to posting a reserve price that applies to the single live demand source. Here the reserve price would be informed by expected demand from static sources and the optimal auction would sell to the static demand source with the highest Value CPM when the live demand source failed to exceed the reserve price.[129] Dynamic Allocation

[125] By live sources, I mean the demand sources that hold live ad auctions, such as exchanges.
[126] See Section III.B.1 for further discussion on the suboptimality of the waterfall format.
[127] Roger B. Myerson. "Optimal Auction Design." *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73.
[128] In the independent private values model. Optimal auctions in the (not necessarily independent) private values, and especially interdependent private values models are significantly more complicated. The optimal format in these more complex settings still solicits bids from all sources simultaneously, and sequential formats are suboptimal.
[129] When restricted to treating the single live demand source as a waterfall line item, several problems might arise. First, if the live demand source is anywhere but first in the waterfall, the impression might be sold before the live source is queried, who might have returned a higher bid. Second, if the live demand source is somehow first in the waterfall but with a low reserve price, it might win the impression despite returning a bid below the best static line item. Together, these identify that the waterfall process with one live demand source can only be optimal if the live

matches this general framework, but by default sets a suboptimal reserve on AdX (the single live demand source). In particular, the maximum Value CPM is certainly not the optimal reserve. Optimal pricing always sets a reserve strictly greater than the opportunity cost (when the opportunity cost is static and known).[130, 131, 132]

112. Suboptimality notwithstanding, during a period when AdX was the only live demand source, Dynamic Allocation might lead to an increase in revenue for publishers in comparison to not calling live sources at all. When all other demand sources are static, Dynamic Allocation simply gives the publisher a shot at additional revenue (even if that shot is taken sub-optimally by default). Indeed, internal Google documents introducing Dynamic Allocation present the ability to call a live demand source at all as a primary benefit.[133] Later on, however, other live demand sources (such as competing exchanges) arose, and as a result, this argument for the benefit of Dynamic Allocation is no longer valid. That is, the primary arguments supporting Dynamic Allocation with only static line items as potentially leading to higher revenue for publishers do not at all apply when line items are live. I now analyze the impact of Dynamic Allocation when other line items are themselves live.

2) Dynamic Allocation with live demand sources

113. In the years after the implementation of Dynamic Allocation, publishers gained the ability to integrate multiple live demand sources, such as exchanges, to DFP.[134] With the inclusion of live demand sources, Dynamic Allocation changes the waterfall process in the following manner:

---

source goes first with a reserve exceeding the Value CPM of the best static line item. Dynamic Allocation with Static Line Items allows a sophisticated publisher to satisfy this property, although by default sets a suboptimal (too low) reserve equal to the maximum Value CPM.

[130] A second source of suboptimality is that static line items are measured in 'Value CPM' which are set by the publishers themselves and might be unequal to the CPM (in which AdX is measured), if the ads are of high/low quality, or if discounts were applied. *See* Google. "Value CPM." Accessed on May 31, 2024. https://web.archive.org/web/20221202071803/https:/support.google.com/admanager/answer/177222?hl=en Whereas AdX competes directly with CPM, so the comparison of CPM from AdX to Value CPM from static line items is not quite right. This can again be mitigated by a sophisticated publisher.

[131] Because there is only one live demand source, the three models (independent private values, private values, and interdependent values) are identical (because there is no uncertainty about the static bids). Therefore, Myerson's optimal auction is truly the optimal mechanism with one live demand source and several static demand sources.

[132] As noted, publishers could increase the reserve set on AdX beyond the maximum Value CPM, which would let them recover the optimal format in this special case with only one live demand source, but the default behavior of Dynamic Allocation sets the maximum Value CPM as AdX's reserve.

[133] GOOG-NE-03597611 at -13. December 15, 2011. "Mysteries of Dynamic Allocation." ("Dynamic Allocation […] maximizes publishers' yield […] By serving AdX […] whenever they offer more than the competing booked ad networks (real-time competition).")

[134] *See* Interactive Advertising Bureau. "OpenRTB." Accessed on June 4, 2024. https://web.archive.org/web/20240326073202/https://www.iab.com/guidelines/openrtb/ (OpenRTB protocol predates header bidding and gave the publishers the ability to sell impressions through live demand sources.)

a. First, DFP processes the high priority line items (such as direct deals). If any high priority line items succeed, the impression is sold, and the waterfall terminates without continuing to subsequent steps.

b. Every lower priority static line item, including AdX, has both a price floor and a Value CPM. Next, DFP computes the highest Value CPM among all low priority line items that satisfy the targeting criteria for this impression, and stores this as a reserve price r.

   i. If the relevant line item is static, the Value CPM is equal to the (value-adjusted)[135] CPM earned by the publisher for selecting that line item.[136]

   ii. If the relevant line item is an exchange, the Value CPM is equal to the (value-adjusted) average historical CPM paid by this exchange to this publisher for past similar impressions.[137, 138]

   iii. If the relevant line item is a header bidding bid,[139] the Value CPM is equal to the header bid.[140] Note that if other exchanges participate primarily via header bidding, this makes the maximum value CPM equal to the clearing price of the header bidding auction.[141, 142]

---

135 I use the parenthetical (value-adjusted) to note that publishers may care about more than just the revenue earned (for example due to ad quality, a discount offered, *etc.*).
136 A sophisticated publisher could ignore Google's suggested formulas and set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.
137 As noted above, this is how publishers typically set Value CPMs, although they are free to deviate from this default behavior. *See* GOOG-AT-MDL-008842393 at -96. August 4, 2023. "Declaration of Nitish Korula." ("Up until at least December 2021, publishers could set the CPM for their booked static remnant line items (also referred to as "Value CPMs").")
138 A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.
139 DFP may not know that the relevant line item is a header bidding bid, only that it is a remnant line item. This does not change any conclusions I make.
140 Notice that the header bidding bid is neither static nor a historical average, it is the winning bid from the header bidding auction. See Section III.B.2 for more details on how header bidding auctions are conducted.
141 Publishers had the ability to increase the clearing price passed on from their header bidding setup to DFP, such as with a multiplier, or an added value. *See* Asmaâ Bentahar. "Bid Adjustments Simplified: Run Fair Auctions with no Hassle" (May 2, 2021). Accessed on May 31, 2024. https://web.archive.org/web/20231202021004/https:/www.pubstack.io/topics/bid-adjustments-simplified (The webpage explains how to implement "bid adjustments" under the section "What are my current Bid Adjustments, and how do I update them?")
142 A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.

c. Next, DFP calls AdX with reserve price equal to the maximum of *r* and AdX's price floor. If AdX succeeds, the impression is sold through AdX, and the waterfall terminates without continuing to subsequent steps.

d. Finally, if AdX fails to return a clearing price higher than this reserve, the ad server visits low priority line items one at a time in decreasing order of Value CPM. When a static line item is visited, it immediately clears the impression. When a header bidding line item is visited, it immediately clears the impression. When an exchange line item is visited, it runs a live auction for the impression, which may or may not win, so the waterfall might continue after visiting an exchange line item if the impression does not clear.

114. In particular, the standard static waterfall auction, without applying Dynamic Allocation, proceeds by executing steps (a) and (d) above, in which all demand sources, including AdX, are processed in order of their Value CPM (which is static). Dynamic Allocation with Static Line Items is a special case of Dynamic Allocation without exchange or header bidding line items.

115. An internal Google deck[143] details the explanation laid out above, an excerpt from which can be seen in Figure 21. It notes that the other demand sources compete based on historical CPM, whereas AdX competes with real-time bids.

[143] GOOG-DOJ-27769247. September 2, 2016. "Header Bidding and FAN."

Figure 21: An excerpt from an internal Google deck that explains how Dynamic Allocation with live demand sources works[144]

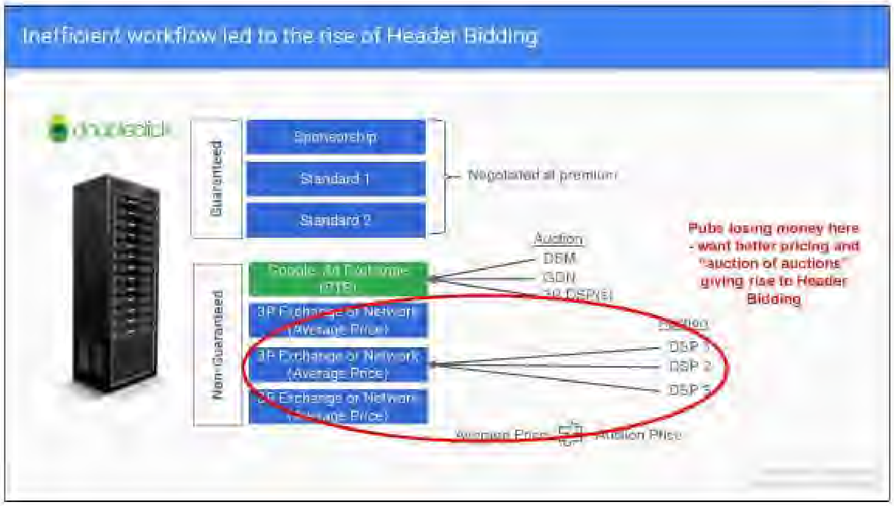

The way DFP works, AdX gets access to all the non-guaranteed inventory and price it on a real time basis.

Whereas other exchanges or networks, only compete on the basis of historical average price - which means they don't get to see all the inventory and make a real time bid, leading to lower yields for publishers.

116. To illustrate how Dynamic Allocation with live demand sources work, imagine an impression arrives for a user over the age of 25 who likes running. Before DFP executes, header bidding[145] solicits bids from the exchanges integrated into the publisher's header bidding setup. The highest bidder that exceeds their personalized reserve is Nike for $4 through OpenX, and this is entered as a line item in DFP. DFP then begins the waterfall and notices that this satisfies the targeting criteria for a direct deal with Altra, and decides to display Altra's ad. The waterfall (even with Dynamic Allocation) terminates here. This example is illustrated in Figure 22 below.

[144] GOOG-DOJ-27769247 at -68. September 2, 2016. "Header Bidding and FAN."
[145] That is, if the publisher set up header bidding auction on their website. For this example, I assume the publisher did.

Figure 22: An impression is allocated to an Altra direct deal under Dynamic Allocation because it satisfies the targeting criteria

117. Another impression arrives for a user over the age of 25 who likes expensive handbags. Before DFP executes, the header bidding auction solicits bids from the header bidding exchanges. The highest bidder that exceeds their personalized reserve is Prada for $4.50 through OpenX, and this is entered as a line item in DFP, which then begins the waterfall. DFP does not have a direct deal that meets this targeting criteria, and so it moves on to lower priority line items and observes that the highest Value CPM option is Prada's header bidding winning bid of $4.50. DFP then calls AdX with a reserve of $4.50. AdX does not find a buyer above $4.50. The waterfall then concludes by selling the impression to Prada for $4.50. This example is illustrated in Figure 23 below.

Figure 23: An impression is allocated to the header bidding winner because AdX fails to clear the Dynamic Allocation reserve

118. A final impression arrives for a user over the age of 25 who likes expensive handbags, and with further fine-grained cookies noting that the user has recently visited several handbag-shopping websites, lives in a wealthy neighborhood. Before DFP executes, the header bidding auction solicits bids from the exchanges integrated into the publisher's header bidding setup. The highest bidder that exceeds their personalized reserve is Prada for $4.50 through OpenX, and this is entered as a line item in DFP, which then begins the waterfall. DFP does not have a direct deal that meets the coarse targeting criteria, and so moves on to lower priority line items and observes that the highest Value CPM option is Prada's header bidding winning bid of $4.50. DFP then calls AdX with a reserve of $4.50 and the auction concludes with Gucci winning at a clearing price of $6. Gucci wins the impression through AdX, paying $6.[146] This example is illustrated in Figure 24 below.

[146] This example abstracts away from the ad buying tools and ad exchange fees for clarity.

Figure 24: Gucci wins an impression through AdX since it clears the Dynamic Allocation reserve

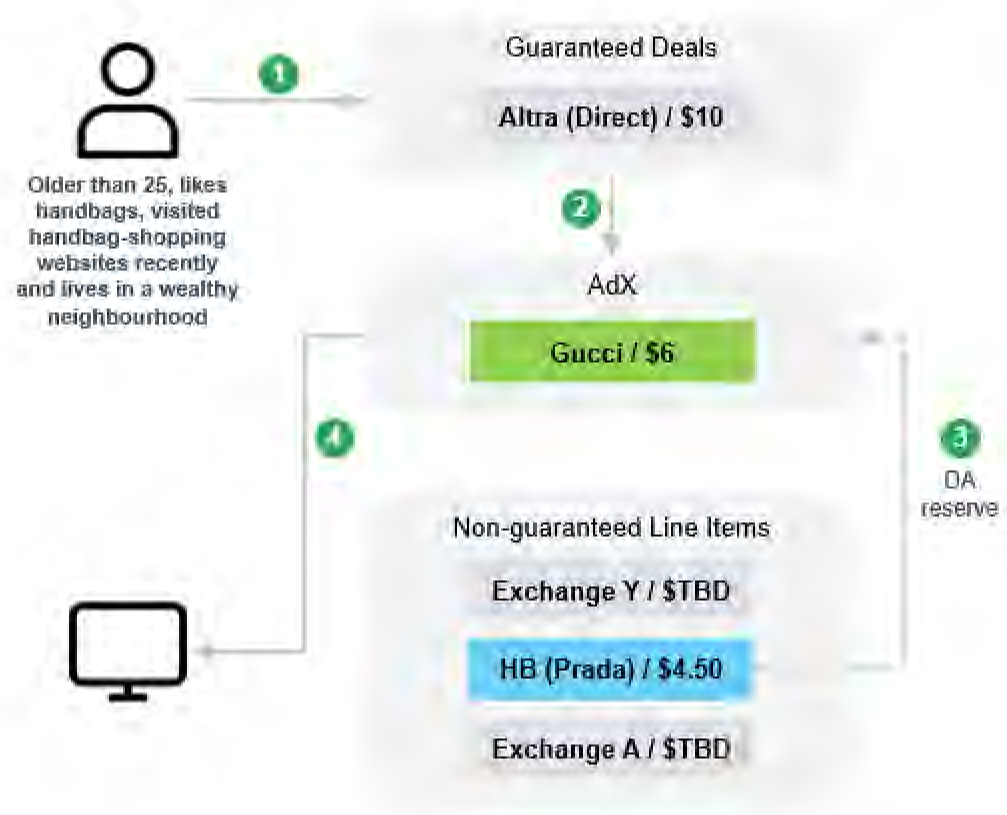

119. If all exchanges, including AdX, participated in header bidding, as opposed to the Dynamic Allocation process, this would create a standard first-price auction with personalized reserves.[147] When one exchange is singled out for default Dynamic Allocation, that exchange is clearly advantaged.[148] In essence, the auction format in this hypothetical is now: (a) all exchanges except for AdX submit a bid in a sealed bid format without learning any others' bids, (b) AdX sees the current highest bid, and then submits its own, (c) the highest bid wins and pays their bid.[149] Hence, Dynamic Allocation allows AdX (and only AdX) to learn others' bids in a first-price auction format, and as a result, Dynamic Allocation creates information asymmetries that favor Google's AdX.[150] This advantage is often referred to as AdX's *Last Look advantage*. I further overview this concept in Section V.C.

[147] It is potentially a bit more complicated than this, only because the bidders in this auction are exchanges rather than direct buyers, but aside from this complication it would just be a first-price auction with personalized reserves.
[148] If a sophisticated publisher uses Dynamic Allocation's flexibility to boost the AdX reserve, Dynamic Allocation still creates an information asymmetry, although the ultimate impact on AdX is indeterminate. I discuss this in more detail when I discuss Last Look in Section V.C.
[149] In reality, it is more complicated than this, since some line items are static and some are "live", but a static bid is still a bid.
[150] See Section II.C for further discussion on how learning others' bids can create advantages in a first-price auction.

B. The Impact of Dynamic Allocation

120. Based on the above articulation of how Dynamic Allocation works,[151] I can draw conclusions on its effects. In my opinion, Dynamic Allocation led to higher win rate and higher revenue for AdX as well as lower win rate and lower revenue for non-Google exchanges.[152] Additionally, if AdX typically transacts ads of lower quality than non-Google exchanges, Dynamic Allocation also led to an increase in the display of lower quality ads.[153] This is my opinion in aggregate, accounting for the possibility that some publishers chose to use default options while others chose to cleverly set Value CPMs to optimize revenue, and accounting for periods both when other exchanges participated via the waterfall and when other exchanges participated via header bidding. In the subsections below, I draw more precise conclusions for each of these settings.

1) Impact of Dynamic Allocation when other exchanges participate in the waterfall

121. I first compare AdX to exchanges that participate in the waterfall. In this case, Dynamic Allocation has the following effect as compared to no Dynamic Allocation: (a) AdX is always visited first in the waterfall, and (b) AdX's reserve may be increased. I now draw conclusions of these effects.

122. In comparison to exchanges that participate in the waterfall, under Dynamic Allocation, and no matter how a publisher sets Value CPMs, AdX is the only exchange that always has the opportunity to submit a bid on any inventory that is not sold through high priority line items.[154] More specifically, if AdX returns a bid that exceeds its reserve price set by Dynamic Allocation, then no other exchange has the opportunity to submit a bid because the waterfall stops, so therefore, AdX is the only exchange that always has this opportunity.[155]

[151] For the entirety of this section, when I say 'Dynamic Allocation,' I refer to 'Dynamic Allocation with Live Demand Sources' and not 'Dynamic Allocation with Static Line Items' (which I introduce both for historical context and to develop intuition as a special case).
[152] In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.
[153] In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.
[154] This is only true when other exchanges are in the waterfall. If other exchanges can access header bidding, they all have an opportunity to bid on each impression.
[155] I again note that AdX would be more likely to return a bid that exceeds its Dynamic Allocation reserve price due to conduct such as Dynamic Revenue Sharing that increases the likelihood that AdX clears its publisher-set reserve price.

a) Dynamic Allocation enables AdX to transact more high-value impressions

123. Let me next consider "high-value" impressions, which have relatively high value given the fine-grained targeting data available to live bidders, but cannot be recognized as such merely on the basis of coarse targeting data used to set static reserves.[156] When other exchanges primarily participate via the waterfall, Dynamic Allocation, no matter how a publisher sets Value CPMs, would lead to AdX winning an even greater volume of high-value impressions, and increased revenues from these impressions under Dynamic Allocation compared to no Dynamic Allocation.[157] This is because AdX submits a live bid based on fine-grained targeting data, whereas its static reserve price is set based only on coarse targeting data (no matter how publishers set Value CPMs, they are still static). If AdX is winning a higher volume of these high-value impressions, then other demand sources necessarily win a lower volume of these high-value impressions. This is one instance where the benefit of going first clearly outweighs the cost of a potentially higher reserve.

124. The impact on "typical-value" impressions is less clear-cut. On one hand, under Dynamic Allocation, AdX has the ability to make *all* impressions available first, meaning that AdX will always get an opportunity to solicit bids. On the other, Dynamic Allocation results in a higher reserve on AdX than without Dynamic Allocation, and so AdX's solicited bids are less likely to clear its reserve. For impressions of higher-than-average value, the prior reasoning still holds; the cost of facing a higher reserve is still minimal (because it is static), whereas the benefit of a first bite is significant (because live bids are likely to be higher than any of the static reserves).[158] For impressions of lower-than-average value, the cost of facing a higher reserve could outweigh the benefit of going first.

2) Impact of Dynamic Allocation when other exchanges participate in header bidding

125. I now compare AdX to exchanges that participate via header bidding. In this case, via Dynamic Allocation, AdX learns information about bids relayed by other exchanges.

---

[156] To have an example in mind, imagine a male over the age of 25 who likes running, *and has visited ten different shoe company's websites in the last hour*. Static reserves based on coarse targeting information would treat this impression like any other male over the age of 25 who likes running, while live bids will know that this is an unusually high-yield impression.

[157] In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.

[158] In my opinion, conduct such as Dynamic Revenue Sharing (which causes AdX to clear its publisher-set price floor more often) would lower the bar for an impression to be considered "higher-than-average" value.

a) Dynamic Allocation enabled AdX to learn the header bidding clearing price

126. In comparison to exchanges that participate via header bidding, under Dynamic Allocation, AdX is the only exchange that learns information about others' bids and passes it on to its bidders. In particular, all exchanges that participate in header bidding must relay bids that are made without any information regarding other exchanges' bids. If the publisher uses default options, AdX can instead relay bids while knowing the maximum bid returned by all header bidding exchanges. I previously noted that this is referred to as AdX's Last Look advantage and constitutes a significant advantage in an auction (see Section V.C).[159]

b) Dynamic Allocation enabled AdX to win impressions by bidding one cent above the header bidding clearing price

127. To illustrate why the information on other participant's bids is useful in a first-price auction, imagine that Bidder One has a value of \$8 and Bidder Two has a value of \$10, and that they participate in a sealed bid first-price auction. Bidder One would surely not submit a bid of \$8, as this guarantees them no gain (either they lose and gain nothing, or they win and pay \$8, which is again leads to no gain). Bidder One does not know what bid Bidder Two will submit, and so does not know how much to shade their bid. Still, Bidder One may form a belief about Bidder Two's behavior, and compute that their optimal bid, given the information they have, is \$4. Bidder Two similarly will not submit a bid of \$10, similarly does not how much to shade their bid, and perhaps computes that their optimal bid in response to whatever belief they form about Bidder One's behavior is to shade their bid to \$5. In this case, Bidder Two wins, pays \$5, and enjoys a utility of \$10-\$5 = \$5. Instead, imagine that Bidder One learns Bidder Two's bid before submitting their own. Bidder Two still will surely not submit a bid of \$10. In particular, while they do not know Bidder One's value, Bidder Two guesses that Bidder One will likely outbid them by a penny if Bidder One's value exceeds Bidder Two's bid and lose otherwise. Uncertainty over Bidder One's value may still result in Bidder Two optimally shading their bid to \$5. In this case, Bidder One would then submit a bid of \$5.01 to win the auction, paying \$5.01 and enjoying a utility of \$8−\$5.01 = \$2.99. This example demonstrates the mechanics by which a bidder would win more often and gain by seeing others' bids in a first-price auction. I also provide a more detailed example in Appendix D to more precisely convey this principle.

[159] If a sophisticated publisher instead cleverly sets Value CPMs as a function of header bids, then AdX might still infer information about the maximum header bid, and in particular certainly knows that the maximum header bid lies below its reserve.

128. In the interest of tractable intuition, the above example is simplified (there is no reserve, bidders bid directly without an exchange, etc.). The same principles extend to Dynamic Allocation where the publisher uses default options.

a. The fact that there is no reserve simplifies the math but is not material to the conclusions. As long as the publisher uses default options in Dynamic Allocation, Bidder One in this example can still win by bidding a penny above Bidder Two when they both exceed the reserve.

b. The fact that the bidders directly bid in the auction instead of going through exchanges is immaterial if those exchanges themselves run first-price auctions.[160, 161, 162] Indeed, bidders who bid in Exchange One still learn the maximum bid from Exchange Two, which still better informs their optimal bid-shading.

c. If bidders instead bid directly through exchanges that run second-price auctions, the benefit is even starker. A second-price auction executed by Exchange One is truthful, even given the existence of Exchange Two. More importantly, given that Exchange Two's second highest bid sets the reserve for Exchange One, Exchange One will win whenever their highest value exceeds Exchange Two's second highest bid.[163] On the other hand, Exchange Two wins only when its second highest bid exceeds the highest bid of Exchange One.

[160] Many exchanges, such as AppNexus, Index Exchange and OpenX gradually switched to the first-price auction format. *See* Sarah Sluis. "Big Changes Coming To Auctions, As Exchanges Roll The Dice On First-Price" (September 5, 2017). Accessed on May 31, 2024. https://web.archive.org/web/20220712083559/https:/www.adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price/

[161] A Google engineer stated that "Many exchanges began to move from second-price to first-price auctions in the mid-to-late-2010s." GOOG-AT-MDL-008842393 at -95. August 4, 2023. "Declaration of Nitish Korula."

[162] AdX later switched to the first-price auction format as well, in 2019. *See* Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20240122142404/https:/blog.google/products/admanager/update-first-price-auctions-google-ad-manager/

[163] When an exchange uses a second-price auction, this auction itself is truthful. However, without a Last Look advantage, the clearing price of this exchange is then entered into a subsequent first-price auction via header bidding. Therefore, the advertiser has to determine not only their own bid, but ways in which they can affect the clearing price (for example, they might want their clearing price to be higher so that their bid wins in the header bidding auction). While their own winning bid cannot affect the clearing price, a very sophisticated advertiser could create another account and "second-price themselves" in order to increase the clearing price they pay and improve their chances of winning the subsequent first-price auction. All of this is to say that bidding in a truthful second-price auction when that auction is run by an exchange that needs to submit the winning bid to a subsequent first-price auction is complicated. On the other hand, bidding in a second-price auction in an exchange with Last Look advantage is still truthful. As long as a bidder submits a bid that exceeds the maximum bids returned by all other exchanges, and all other advertisers who participate in AdX, that bidder will win (and the price will be their minimum bid to win). Therefore, an exchange running a second-price auction with Last Look advantage provides an advantage to its bidders as compared to an exchange running a second-price auction without a Last Look advantage.

129. I therefore conclude that, when exchanges participate primarily via header bidding, AdX would win more impressions under Dynamic Allocation with default options as compared to no Dynamic Allocation.[164] Because AdX earns revenue via take-rates on cleared transactions, this would cause AdX's revenue to increase. Because this conclusion applies to any impression, it applies to high-value impressions as well. By the same logic, non-Google exchanges would necessarily clear fewer transactions (including high-value transactions) and earn less revenues.[165]

3) Impact of Dynamic Allocation on ad quality

130. If it is the case that AdX typically transacts ads that are of lower quality compared to non-Google Exchanges, then an increased win rate for AdX would result in lower quality ads for the publisher.[166] Based on my above conclusions, Dynamic Allocation, no matter how publishers set Value CPMs, would result in lower quality ads displayed on high-value impressions when exchanges participate primarily through the waterfall. Similarly, Dynamic Allocation with default options would result in lower quality ads displayed on all impressions when exchanges participate primarily through header bidding. This follows because Dynamic Allocation results in AdX winning additional impressions, and displaying ads of whatever quality it tends to display.

C. Enhanced Dynamic Allocation

131. Google later amended Dynamic Allocation to include guaranteed line items as well. This is referred to as **Enhanced Dynamic Allocation** and it changes the waterfall process in the following manner:

a. First, DFP calculates a "temporary CPM"[167] for all high priority line items according to the campaign goals associated with those line items. For example, if the delivery is behind for a specific high priority line item, the temporary CPM associated with

[164] In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.
[165] The impact of Dynamic Allocation with sophisticated publishers who cleverly set Value CPMs is less clear-cut. On one hand, if sophisticated publishers only slightly inflate the Value CPM of the winning header bid, then the above conclusions continue to hold for exactly the same reasons. On the other hand, if sophisticated publishers significantly inflate the Value CPM of the winning header bid due to Dynamic Allocation and would not have set such an inflated reserve on AdX in absence of Dynamic Allocation, then the cost of this inflated reserve might outweigh the benefits highlighted above.
[166] In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.
[167] GOOG-AT-MDL-008842393 at -99. August 4, 2023. "Declaration of Nitish Korula." ("Up to at least December 2021, with Enhanced Dynamic Allocation, the ad server calculated what is known as a temporary CPM for a guaranteed deal.")

that line item goes up.[168] DFP stores the highest temporary CPM as a reserve price *R*.

b. Next, DFP computes the highest Value CPM among all low priority line items that satisfy the targeting criteria for this impression, and stores this as a reserve price *r*.

   i. If the relevant line item is static, the Value CPM is equal to the (value-adjusted) CPM earned by the publisher for selecting that line item.[169]

   ii. If the relevant line item is an exchange, the Value CPM is equal to the (value-adjusted) average historical CPM paid by this exchange to this publisher for past similar impressions.[170]

   iii. If the relevant line item is a header bidding bid, the Value CPM is equal to the header bid. Note that if other exchanges participate primarily via header bidding, this makes the maximum value CPM equal to the clearing price of the header bidding auction.[171]

c. Next, DFP calls AdX with reserve price equal to the maximum of *r, R,* and AdX's price floor. If AdX succeeds, the impression is sold through AdX, and the waterfall terminates without continuing to subsequent steps.

d. Finally, if AdX fails to return a clearing price higher than this reserve,[172]

[168] A Google engineer explained in a declaration that [redacted] GOOG-AT-MDL-008842393 at -99. August 4, 2023. "Declaration of Nitish Korula."

[169] A sophisticated publisher could ignore Google's suggested formulas and set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.

[170] A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.

[171] A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.

[172] GOOG-AT-MDL-008842393 at -00. August 4, 2023. "Declaration of Nitish Korula." ("if the highest effective AdX bid could beat both the EDA price and the price of the remnant line item that was selected as a candidate for the impression, then the ad associated with that AdX bid would win. If not, the guaranteed or remnant line item would win.") This suggests that the winner in this case is determined by a comparison between r and R and does not

i. If $R > r$, the impression is sold to the high priority line item with the highest temporary CPM (equal to R).

ii. Otherwise, the impression is offered to the low priority line item with the highest Value CPM (equal to r). If this line item is static or a header bid, the impression is immediately cleared. If this line item is an exchange, the exchange runs a live auction with its assigned price floor, and the impression may or may not clear.

132. The key difference between Dynamic Allocation and Enhanced Dynamic Allocation is that high priority line items no longer retain their absolute priority over AdX, which is demonstrated in step c above. Step d describes what happens if AdX fails to clear its reserve.

133. The following excerpt in Figure 25 from an internal Google document on Google auction adjustments[173] outlines how the reserve price is determined by Enhanced Dynamic Allocation. It corroborates the explanation I presented at the beginning of this subsection.

explicitly state what happens in case the remnant line item is an exchange that does not clear the impression. The Google documentation I have access to does not explicitly clarify this aspect, but exactly how this case resolves is immaterial to my conclusions.

[173] GOOG-NE-13203009**.** “DRX Global Optimization of DRS, RPO, and EDA.”

Figure 25: An excerpt from an internal Google slide deck on Enhanced Dynamic Allocation, outlining that the dynamic reserve price is a calculation based on both the delivery of the guaranteed line items as well as the maximum valued non-guaranteed line item[174]

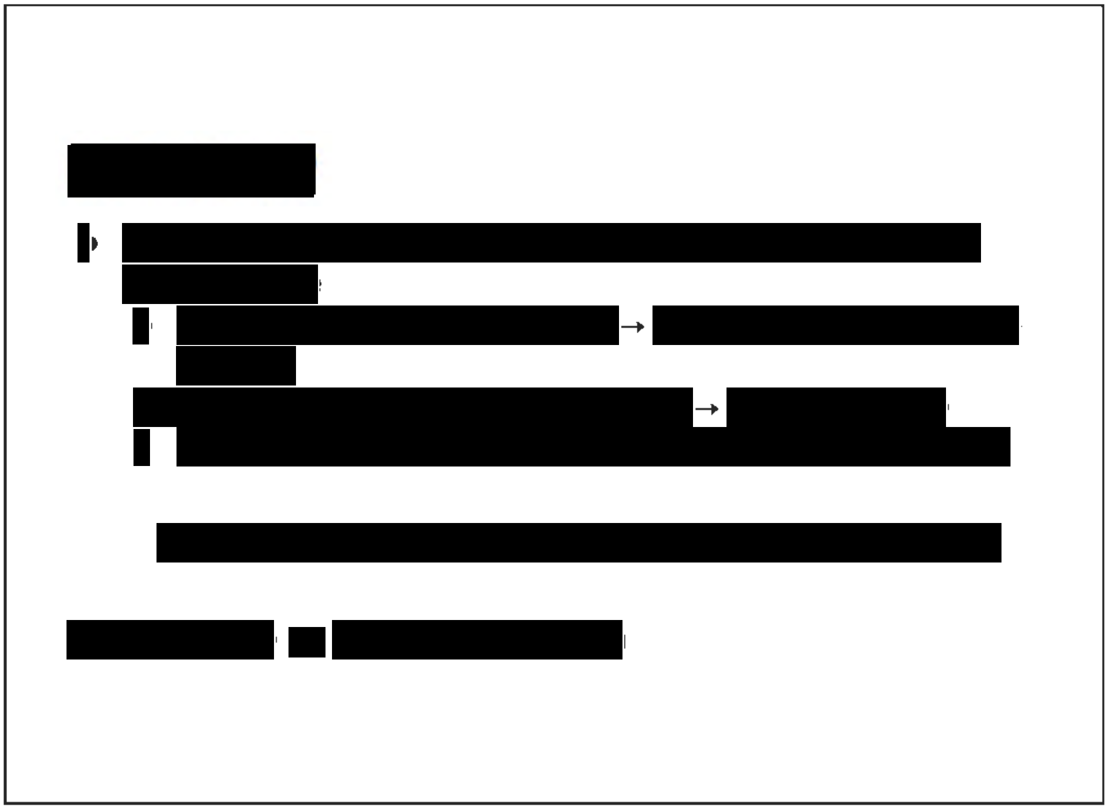
[redacted]

134. To illustrate how Enhanced Dynamic Allocation impacts the auction process, imagine an impression arrives for a user over the age of 25 who likes running. Before DFP executes, the header bidding auction solicits bids from exchanges they integrated to their header bidding system. The highest bidder that exceeds their personalized reserve is Nike for $4 through OpenX, and this is entered as a line item in DFP, which then begins the waterfall, and notices that this satisfies the targeting criteria for a direct deal with Altra, with a Value CPM of $10, which is the highest Value CPM among all line items. DFP then calls AdX with a reserve price of $10,[175] but AdX does not find a buyer above $10. The waterfall then concludes by using the impression to fulfill Altra's direct deal. This example is illustrated in Figure 26 below.

[174] GOOG-NE-13203009 at -16. "DRX Global Optimization of DRS, RPO, and EDA."
[175] This example assumes that the guaranteed line item price directly becomes the Enhanced Dynamic Allocation without the "temporary CPM" adjustment for the sake of clarity.

Figure 26: An impression is allocated to an Altra direct deal because AdX fails to clear the Enhanced Dynamic Allocation reserve[176]

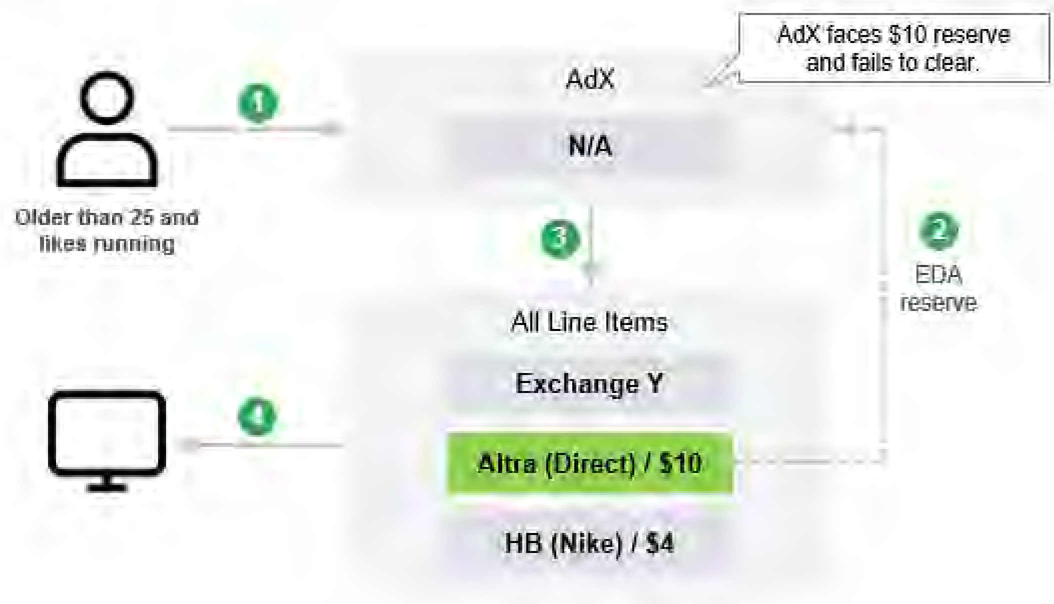

135. Another impression arrives for a user over the age of 25 who likes running, and with further fine-grained cookies noting that the user has recently visited several shoe-shopping websites, lives in a wealthy neighborhood, and competes in races. Before DFP executes, the header bidding auction solicits bids from the exchanges they integrated to their header bidding system. The highest bidder that exceeds their personalized reserve is Nike for $8 through OpenX, and this is entered as a line item in DFP, which then begins the waterfall, and notices that this satisfies the targeting criteria for a direct deal with Altra, with a Value CPM of $10, which is the highest Value CPM among all line items. DFP then calls AdX with a reserve of $10,[177] whose auction concludes with Reebok winning at a clearing price of $15. Reebok wins the impression through AdX, paying $15.[178] This example is illustrated in Figure 27 below.

[176] This example assumes that the guaranteed line item price directly becomes the Enhanced Dynamic Allocation without the "temporary CPM" adjustment for the sake of clarity.
[177] This example assumes that the guaranteed line item price directly becomes the Enhanced Dynamic Allocation without the "temporary CPM" adjustment for the sake of clarity.
[178] This example abstracts away from the ad buying tools and ad exchange fees for clarity.

Figure 27: Reebok wins an impression through AdX because it is able to clear the Enhance Dynamic Allocation reserve set by the Altra direct deal[179]

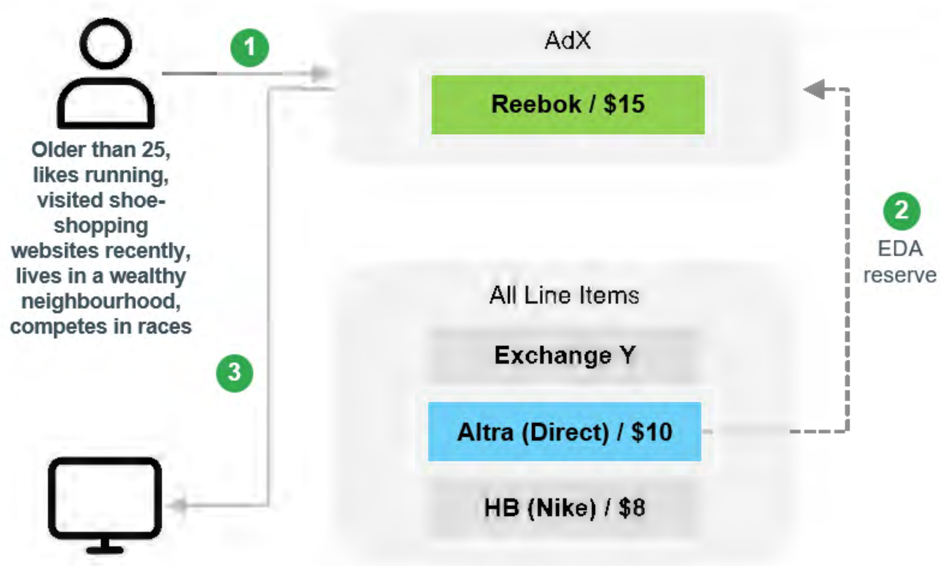

136. These examples show that Enhanced Dynamic Allocation gives AdX the opportunity to bid against high priority line items such as direct deals, too. The last example also highlights one reason that AdX might possibly outbid high priority line items which typically have a much higher CPM compared to remnant line items,[180] that the fine-grained targeting criteria may alert live demand sources of a particularly high value impression and bid higher than the direct deal price that is typically based on coarser targeting criteria.

D. The Impact of Enhanced Dynamic Allocation

137. With the above articulation of how Enhanced Dynamic Allocation works, I now draw conclusions on its effects. In my opinion, Enhanced Dynamic Allocation likely led to an increase in win rate and increase in revenue for AdX and reduced the value of direct deals for advertisers, which would in turn decrease the revenue earned by publishers via direct deals. [181] I explain all these conclusions below.

---

[179] This example assumes that the guaranteed line item price directly becomes the Enhanced Dynamic Allocation without the "temporary CPM" adjustment for the sake of clarity.

[180] Google states that "On average, direct ads sell for two to four times higher than programmatic ads. Direct ads average $10-20 CPMs (cost per thousand impressions), while programmatic ads average $1-5 CPMs." Google. "Understand Direct and Programmatic Ad Revenue." Accessed on May 31, 2024. https://web.archive.org/web/20231226200704/https:/newsinitiative.withgoogle.com/resources/trainings/grow-digital-ad-revenue/understand-direct-and-programmatic-ad-revenue/

[181] In my opinion, the magnitude of AdX's win rate and revenue increase would be larger due to conducts such as Dynamic Revenue Sharing which increase the likelihood that AdX clears its publisher-set reserve. The impact of such conducts on direct deals is less clear-cut, and so I do not opine on the further impact of such conducts on direct deals beyond the immediate impact of Enhanced Dynamic Allocation.

138. Enhanced Dynamic Allocation expands Dynamic Allocation to also include high-priority line items. In particular, the 'Dynamic Allocation Framework' can be interpreted as (a) finding a set of line items, then (b) allowing AdX to be considered before any of those line items,[182] but (c) setting the reserve price of AdX to be equal to the maximum Value/Temporary CPM of those line items.[183]

139. Under Enhanced Dynamic Allocation, impressions that otherwise would have been reserved for high priority line items such as direct deals are instead available for AdX's auction. This follows immediately from the definition of Enhanced Dynamic Allocation since it allows AdX to run an auction for all impressions. The Enhanced Dynamic Allocation-generated reserve price might be high, but AdX will always have the opportunity to run an auction for all impressions. Without Enhanced Dynamic Allocation, any impression with a viable high priority line item would not be available to AdX for auction.

140. Furthermore, AdX is the only exchange that unconditionally has this opportunity. Under Enhanced Dynamic Allocation, another exchange *can* have this opportunity, but only if (a) its Value CPM exceeds the highest temporary CPM among high priority line items, and (b) AdX fails to clear its reserve. In particular, (a) suggests a high barrier to this exchange being considered in front of the high priority line item at all,[184] and (b) notes that AdX still gets a first bite, even if the Value CPM of an exchange is high enough to satisfy (a).

141. As Enhanced Dynamic Allocation runs live auctions for every impression, it will likely create a revenue increase for the publishers in the short run.[185, 186] This conclusion is supported by an internal Google slide deck,[187] an excerpt from which is reproduced in Figure 28. It shows that Enhanced Dynamic Allocation led to an increase in publisher revenue from AdX in the first

---

[182] When one of these line items is a header bidding line item, Dynamic Allocation considers AdX before that line item, although of course that line item itself was generated from a live bid in an auction executed before the ad server.
[183] Under default publisher behavior. As previously noted, a sophisticated publisher could further increase AdX's reserve beyond the maximum Value CPM.
[184] Recall that "On average, direct ads sell for two to four times higher than programmatic ads. Direct ads average $10-20 CPMs (cost per thousand impressions), while programmatic ads average $1-5 CPMs." Google. "Understand Direct and Programmatic Ad Revenue." Accessed on May 31, 2024.
https://web.archive.org/web/20231226200704/https:/newsinitiative.withgoogle.com/resources/trainings/grow-digital-ad-revenue/understand-direct-and-programmatic-ad-revenue/
This suggests that temporary CPMs of high priority line items would likely be higher than static Value CPMs of low priority line items.
[185] In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.
[186] Importantly, it is possible that many publishers deprioritize short term revenue, and care more about their revenue in the long run. I elaborate on the long run effects, namely the "cream-skimming effect," below.
[187] GOOG-NE-03872763. "Discussion on improving AdX & AdSense backfill."

quarter of 2014. This also makes sense, as high priority line items are all static. I previously noted in Section IV.A that it would increase revenue to allow a live demand source the option to outbid a static demand source.

<u>Figure 28: An excerpt from an internal Google slide deck plotting the impact of Enhanced Dynamic Allocation on publisher metrics[188]</u>

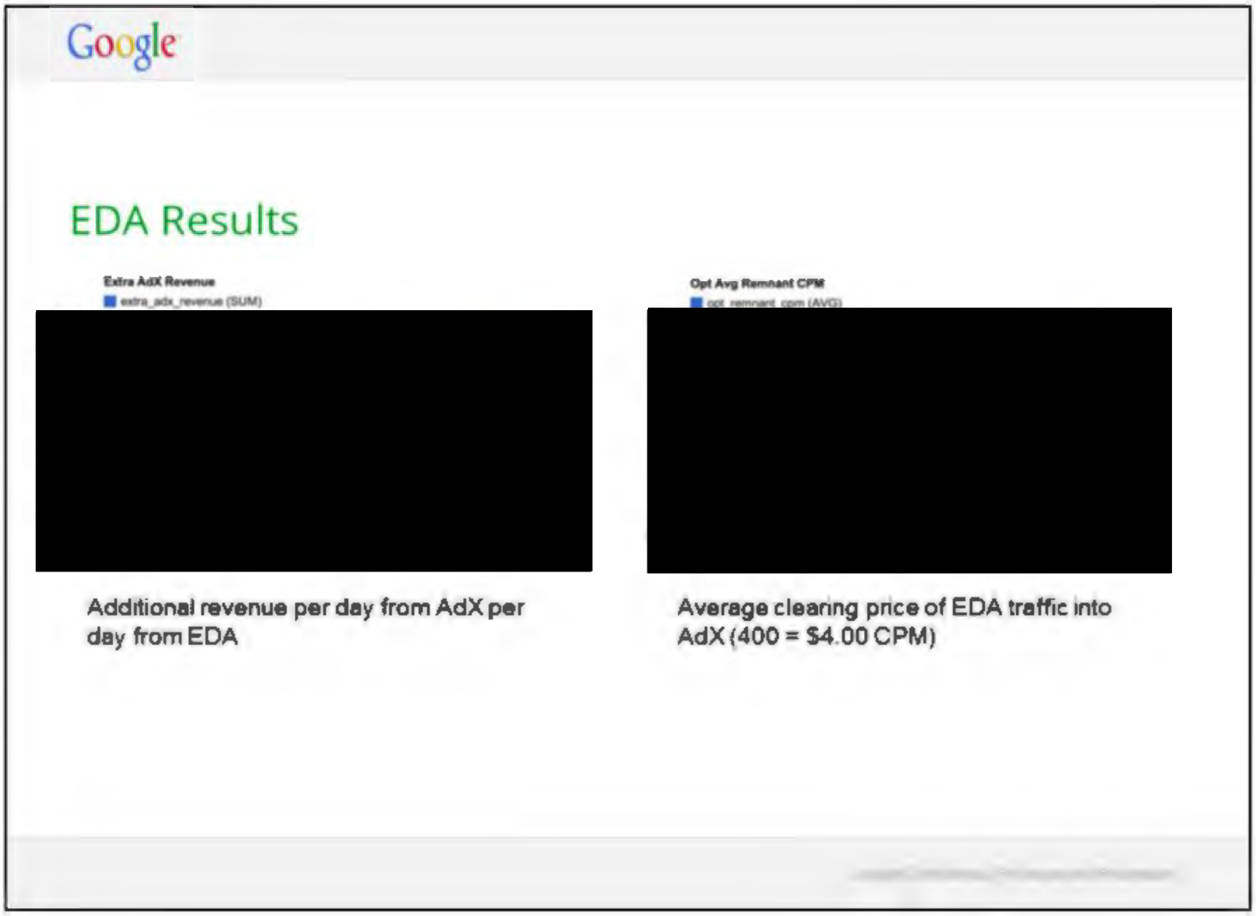

142. Since AdX generates revenue by taking a cut of the clearing prices, the plot above shows that Enhanced Dynamic Allocation increases AdX and Google revenue as well. This can be seen by the upward trends of the plots in the figure above. If impressions that satisfy targeting criteria for direct deals are on average more valuable than impressions that do not,[189] then Enhanced

[188] GOOG-NE-03872763 at -85. "Discussion on improving AdX & AdSense backfill."
[189] Google's online documentation states that "Direct ads average $10-20 CPMs (cost per thousand impressions), while programmatic ads average $1-5 CPMs." *See* Google. "Understand Direct and Programmatic Ad Revenue." Accessed on May 31, 2024.
https://web.archive.org/web/20231226200704/https:/newsinitiative.withgoogle.com/resources/trainings/grow-digital-ad-revenue/understand-direct-and-programmatic-ad-revenue/

Dynamic Allocation results in more valuable transactions being transacted through AdX rather than direct deals. Therefore, Enhanced Dynamic Allocation would not only lead to increased volume and revenues for AdX, but also to a greater volume of valuable impressions being transacted through AdX.[190]

143. This outcome reduces the value of direct deals for advertisers, which would likely decrease the revenue that publishers can expect to earn via direct deals. Specifically, assuming that the value for impressions is refined based on fine-grained targeting data used in live ad auctions after the coarse targeting criteria used in direct deals, live demand sources are better informed on the value of an impression than direct deal partners for that impression. For example, Altra's value for an impression to a man over the age of 25 who likes running might be $10 on average, but sometimes $15 if that user has also visited running shoe websites in the last two hours, and sometimes $5 if that user has not. In the absence of Enhanced Dynamic Allocation and any cream-skimming effect,[191] Altra would enjoy a CPM value of $10 on average for men over the age of 25 who like running, half of whom have also visited running shoe websites in the last two hours and half of whom who have not. But importantly, Altra's direct deal is based on the coarse targeting criteria of men over the age of 25 who like running. With Enhanced Dynamic Allocation, AdX solicits bids based on the full fine-grained targeting criteria. Therefore, if the Temporary CPM of Altra's direct deal is $10, AdX is likely to find a bid exceeding its reserve price for users who have recently visited running shoe websites, and not for those who have not. Therefore, Altra's direct deal would exclusively serve low value impressions to the advertisers, coming from users who have not also visited running shoe websites recently and give a value of $5. Note that this does not limit the publisher's ability to fulfill Altra's direct deal. It simply causes Altra's direct deal to be filled exclusively with low-value impressions rather than a mix of low- and high-value impressions.

144. This cream-skimming effect reduces the value of direct deals with publishers in the long run from the perspective of the advertisers. Hence, it potentially leads to a negative impact on the

[190] In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.

[191] In this example, the "cream" of the impressions to users over the age of 25 who like running are those who have visited running shoe websites in the last two hours. AdX "skims the cream" of these users because it is aware of which users are the cream and which are not when soliciting live bids. Altra, on the other hand, is stuck only with the leftovers (despite the fact that Altra submitted a bid generically on users over the age of 25 who like running, half of whom are the "cream"). This is also an instance of what is called *adverse selection*, because Altra might initially think that they are getting a representative impression among users over the age of 25 (for which their average value is $10), but a better-informed party causes them to receive non-representative impressions for users over the age of 25 (and in particular, the low value ones).

publisher revenue in the long run. There are some means to partially mitigate this impact, but in any auction where Enhanced Dynamic Allocation causes a different outcome than Dynamic Allocation, there is risk of cream-skimming. The key property to consider is that the following two impressions have different values: (a) an impression for an average user who satisfies the coarse targeting criteria, and (b) an impression for an average user who satisfies the coarse targeting criteria and for which AdX does not find a bid exceeding its reserve of r. If r is small, then the values of (a) and (b) might be very far apart. If r is huge, then the values of (a) and (b) might be essentially (or even exactly) the same. Therefore, one mitigation might be to always set temporary CPMs so large that AdX would never find a bid. This is in fact a full mitigation of this effect but results in the same outcome as Dynamic Allocation (and therefore fully nullifies Enhanced Dynamic Allocation). One form of partial mitigation might be to have a reserve price sometimes equal to the true Value CPM of the direct deal, and sometimes so huge that AdX would never find a bid. In the former case, Enhanced Dynamic Allocation results in distinct outcomes from Dynamic Allocation, but the direct deal suffers the full effects of cream-skimming. In the latter, the direct deal does not suffer from cream-skimming, but Enhanced Dynamic Allocation is also just equal to Dynamic Allocation. Another form of partial mitigation might be to have a reserve price higher than the true Value CPM of the direct deal, but not so huge that AdX would never find a bid. In this case, the direct deal suffers some cream-skimming (they only miss out on the exceptionally-high-value impressions), and Enhanced Dynamic Allocation yields only some distinct outcomes from Dynamic Allocation (AdX only has a serious shot at the exceptionally-high-value impressions). These are examples of possible mitigations. For any cream-skimming-mitigation approach, there is a direct tradeoff between how distinct the outcomes are under Enhanced Dynamic Allocation and how much the Direct Deal would be hurt by cream-skimming.

V. CONDUCT ANALYSIS: EXCHANGE BIDDING

145. In this section, I analyze header bidding within the context of auction design, focusing on auction structure and dynamics both prior and subsequent to Google's introduction of Exchange Bidding. My analysis compares header bidding to a waterfall approach (see Section III.B for details) and header bidding to Exchange Bidding, which I discuss in detail below. I find that header bidding improves publisher outcomes relative to the waterfall approach (with and without Dynamic Allocation and Enhanced Dynamic Allocation). I further find that header bidding can generate higher revenue for publishers than can Exchange Bidding.

A. Exchange Bidding

146. I first outline Google's rival product to header bidding. In 2018, Google released **Exchange Bidding**, a technology similar to header bidding, after observing the rise of header bidding technology.[192, 193, 194] Exchange Bidding created a separate auction of auctions into which all exchanges except for AdX submit their bids. AdX would then submit a bid against the prevailing bid from this auction of auctions. Similar to header bidding, Exchange Bidding was a first-price auction.[195] Furthermore, publishers could use both header bidding and Exchange Bidding at the same time.

147. Exchange Bidding augments the waterfall process in the following manner:[196]

a. [redacted]

[redacted]

[192] First the tool was called "Exchange Bidding Dynamic Allocation," later renamed to "Exchange Bidding," then to "Open Bidding" after the co-rollouts of Unified Pricing Rules and the AdX auction format change to the first-price. *See* AdExchanger. "Google's Exchange Bidding Is Now 'Open Bidding'; Market Researchers Slip" (August 27, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20220523024855/https:/www.adexchanger.com/ad-exchange-news/tuesday-27082019/
The internal development code name for the product was "Jedi." *See* GOOG-NE-03995243 at -3. July 25, 2018. "PRD: Unified 1P auction and Pricing rules."
[193] To be more precise, the best product to compare to Exchange Bidding is the header bidding variant called "server-side header bidding," since both handle the auction on a server rather than the user's internet browser. *See* Anthony Vargas. "AdExplainer: Client-Side vs. Server-Side Header Bidding: What's The Difference?" (December 1, 2023). Accessed on May 31, 2024. https://web.archive.org/web/20240314163210/https:/www.adexchanger.com/adexplainer/adexplainer-client-side-vs-server-side-header-bidding-whats-the-difference/ ("With client-side header bidding, the bulk of that processing occurs on the user's device in the web browser itself. With server-side header bidding, the processing happens on a remote server.")
[194] Google's First Am. Resps. and Objs. to Plaintiff's Third Set of Interrogs. (May 24, 2024) at 11.
[195] Internal Google documents state that "EB was our first attempt at running a 1P [first-price] auction; Since other exchanges already have experience with submitting 1P bids into HB wrappers, it was the easiest way to build out the product." GOOG-NE-13494966 at -71. May 2019. "Managing Yield."
[196] GOOG-TEX-00000744. April 26, 2017. "Exchange Bidding (Jedi) Open Beta Sates Readiness Review." (internal Google slide deck that discusses how Exchange Bidding worked during its initial implementation, as well as Google's plans on Exchange Bidding rollout.)

i. [redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

---

[197] A sophisticated publisher could ignore Google's suggested formulas and set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.
[198] A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.
[199] A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default. Recall that some publishers were indeed sophisticated and chose to do this. This is referred to as the "boost." *See, e.g.,* GOOG-TEX-00843142 at -46. September 3, 2019. "First-price bidding Update." ("Why does "boost" exist? The publisher inflates the HB bid before sending it as a floor to AdX. Drives up cost + "fairer" comparison between Google buyer bid and HB bid.")
[200] When integrating an exchange into Exchange Bidding, the publisher has the option to set a personalized price floor.
[201] [redacted]

[redacted]

[redacted]

148. Exchange Bidding can be interpreted as affording some of the advantages provided to AdX to other exchanges,[205] since it extends the Enhanced Dynamic Allocation Last Look advantage to them, provided they pay the Exchange Bidding fee. In fact, the earliest iteration of Exchange Bidding was called "Exchange Bidding in Dynamic Allocation."[206]

149. To illustrate the process of Exchange Bidding, suppose that there is opportunity to serve an impression for a user over the age of 25 who likes running, and with further fine-grained data from cookies noting that the user has recently visited several shoe-shopping websites, lives in a wealthy neighborhood, and competes in races. Before DFP executes, header bidding solicits bids from exchanges. The highest bidder that exceeds their personalized reserve is Nike for $8, and this is entered as a line item in DFP, which then begins the waterfall, and notices that this satisfies the targeting criteria for a direct deal with Altra, with a temporary CPM of $10, which is the highest temporary or Value CPM among all line items. DFP then calls Exchange Bidding and AdX, with a reserve of $20 for AdX, in order to adjust for the fact that the publisher believes AdX ads to be of low quality, and a reserve of $10 for Index Exchange.[207] AdX's auction concludes with no bidder exceeding the reserve, and Index Exchange concludes with Vibram[208]

[202] GOOG-TEX-00105361 at -67. April 28, 2017. "FAN Bidding in to DPI and AdMob."
[203] The bids submitted to the Exchange Bidding auctions are net of this fee, so the amount paid by the advertiser is equal to the bid submitted by the ad buying tool into the exchange plus the ad buying tool fee.
[204] Later, along with the introduction of Unified Pricing Rules, Google switched to what was called "Unified Auction" where all exchanges and all "authorized buyers" participated in a single first-price auction with unified reserves. Since the unified auction was a first-price auction with non-personalized reserves, this meant that AdX effectively switched to first-price auction as well. GOOG-NE-13494966 at -80. May 2019. "Managing Yield."
Under Unified Pricing Rules, publishers cannot set personalized reserve prices on each exchange and must instead set the same reserve price for all exchanges in Exchange Bidding as well as AdX (but this reserve may still exceed *r*). I analyze Unified Pricing Rules in Section VI.
[205] GOOG-DOJ-AT-01809483 at -84. March 2017. "Exchange Bidding in Dynamic Allocation (fka Project Jedi)." ("Exchange Bidding in Dynamic Allocation is a project to allow publishers to invite existing third party exchange partners to bid in real-time alongside AdX in Dynamic Allocation.")
[206] *See* Jonathan Bellack. "Improving yield, speed and control with DoubleClick for Publishers First Look and exchange bidding" (April 13, 2016). Accessed on May 31, 2024. https://web.archive.org/web/20240206070512/https:/blog.google/products/admanager/improving-yield-speed-and-control-with-dfp-first-look-and-exchange-bidding/
[207] Index Exchange is an exchange that operates in the ad exchange market.
[208] Vibram is a company that produces shoes.

winning at a clearing price of $12. Vibram wins the impression through Index Exchange, which pockets 10% and passes on $10.8 to DFP. DFP then takes 5% and passes on $10.26 to the publisher as a payout.[209] This example is illustrated in Figure 29 below.

Figure 29: Vibram wins an impression through Exchange Bidding because it clears the Enhanced Dynamic Allocation reserve set by an Altra direct deal

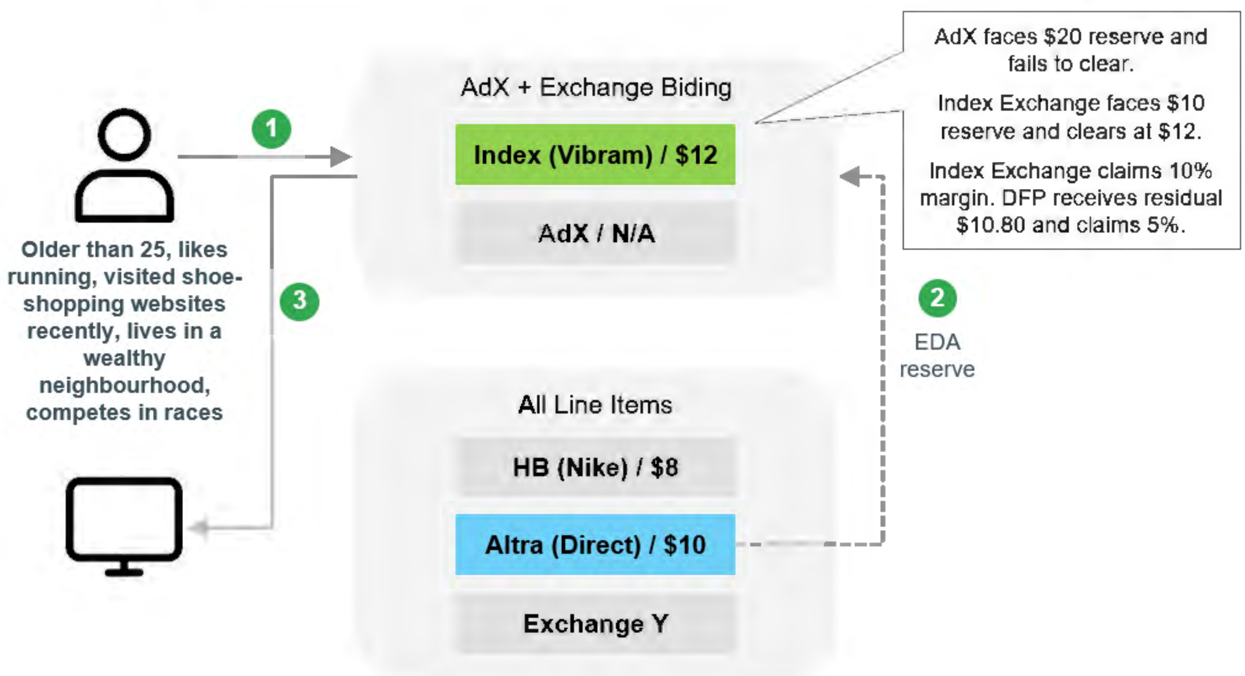

B. Impact of Header Bidding

150. Based on the exchange bidding definition above, along with the definitions of header bidding from Section III.B.2 and Dynamic Allocation and Enhanced Dynamic Allocation from Section IV, I can provide some results regarding the effects of header bidding. In my opinion, header bidding improves publisher revenue and fill rate[210] in comparison to the waterfall process. Furthermore, header bidding's auction mechanics generate higher revenue for publishers than Exchange Bidding's auction mechanics do.

[209] This example abstracts away from details like ad buying tool fees, for clarity.
[210] The fill rate of the publisher is defined as the amount of impressions they were able to sell divided by the total amount of impressions received by the publisher. A higher fill rate implies a lower amount of impressions going unsold, given the total number of impressions.

1) Comparison of header bidding to the waterfalling with and without (Enhanced) Dynamic Allocation

151. Header bidding increased revenue for publishers, in comparison to the waterfall.[211] To illustrate this, consider the bid that might be submitted by an exchange. Under a waterfall approach, this exchange might not be called, in which case the highest bid from this exchange can be interpreted as 0. If called, the highest bid from this exchange would need to meet or exceed their personalized reserve in order to win the impression. In contrast, under header bidding

a. The exchange is certainly called and in order to win must submit a bid that not only exceeds their reserve, but also exceeds the bids of other exchanges. Therefore, this exchange would submit a bid under header bidding that is at least as high as their bid under the waterfall process for all impressions.

b. Furthermore, observe that the winning payment under header bidding is always the maximum bid that exceeds its reserve, whereas the winning payment under waterfalling is the first bid that exceeds its reserve. Therefore, even if individual exchanges' bids were the same in header bidding versus waterfalling, the publisher's revenue would go up because header bidding selects the highest bid instead of the first acceptable bid from the waterfall.

c. Putting both claims together, header bidding would generate bids at least as high as waterfalling, and header bidding generates increased revenue from the same bids compared to waterfalling, and therefore generates increased revenue overall. In particular, the primary source of increased revenue is due to the simultaneous format of header bidding (as compared to the sequential format of the waterfall).

152. This is also confirmed by internal Google documents, one stating that "Pubs are making more money using [header bidding] (20%+ according to pubs) because they are able to get a per-query price from more demand sources."[212]

[211] To be clear, the two scenarios I consider are (a) all exchanges participate in the waterfall with some personalized reserves, and (b) all exchanges participate in header bidding with those same personalized reserves. If the publisher chooses to re-optimize reserves after switching to header bidding, their revenue could only increase further.
[212] GOOG-DOJ-AT-02639830 at -41. April 2016. "Exchange Bidding (aka Jedi) LPS AMERICAS TRAINING."

153. The claims above hold even with the effect of (Enhanced) Dynamic Allocation.[213] I prove this claim in Appendix E. The proof follows from the description of Header Bidding in Section III.B.2 but requires case by case analysis. Again, the primary source of increased revenue is due to the simultaneous versus sequential format (even in light of AdX's Last Look advantage).

154. If header bidding enabled publishers to reach an expanded set of advertisers that they would not have otherwise reached through Google's ad tech stack, one would expect this additional competition to generate greater revenues for publishers.[214] In particular, one would expect existing advertisers to submit higher bids due to increased competition from additional advertisers. As long as existing advertisers' bids do not decrease in response to additional competition,[215] these additional advertisers themselves simply provide additional bids from which the publisher can take the maximum.

155. Google documents support this conclusion as well. An internal Google document[216] states that header bidding is advantageous for publishers because "they aren't getting all of the demand from AdX, so having other exchanges bid for inventory via HB increases Pub yield," showing that header bidding enables publishers to elicit bids from a wider range of advertisers compared to operating solely through the Google ad tech stack.

2) Comparison of header bidding to Exchange Bidding

156. The auction mechanics of header bidding (without Enhanced Dynamic Allocation) would generate increased revenue for publishers, as compared to all exchanges participating in Exchange Bidding. Without (Enhanced) Dynamic Allocation, header bidding is a clean first-price auction with personalized reserves and low take-rate.[217] Exchange Bidding is a clean first-price

---

[213] To be clear, the two scenarios I consider are: (a) non-Google exchanges participate in the waterfall with some personalized reserves and AdX has (Enhanced) Dynamic Allocation, and (b) non-Google exchanges participate in Header Bidding with those same personalized reserves and AdX has (Enhanced) Dynamic Allocation. In particular, for impressions that are eligible for a high priority line item, the two are equivalent. For impressions ineligible for a high priority line item, AdX's reserve in (a) is the maximum of all received header bids and all personalized reserves, and AdX's reserve in (b) is the maximum personalized reserve.

[214] While this may seem like an obvious claim, bidding behavior in first-price auctions is notoriously complex and counterintuitive phenomena are certainly possible. *See* Bernard Lebrun. "Existence of an Equilibrium in First Price Auctions." *Economic Theory* vol. 7, no. 3. 1996. pg. 421–443; Bernard Lebrun. "First Price Auctions in the Asymmetric N Bidder Case." *International Economic Review* vol. 40, no. 1. 1999. pg. 125–142; Bernard Lebrun. "Uniqueness of the equilibrium in first-price auctions." *Games and Economic Behavior* vol. 55, no. 1. 2006. pg. 131-151.

[215] Which would not happen in a second-price auction format, since it is truthful.

[216] GOOG-AT-MDL-001811992. June 2017. "Exchange Bidding / Platform StratOps Meeting."

[217] Client-side header bidding is free through Prebid, an open-source software. *See* Prebid. "Boost Programmatic Advertising Revenue." Accessed on June 3, 2024. https://prebid.org/

auction with personalized reserves (or non-personalized reserves, under UPR) where DFP collects a 5% take-rate on top of the ad exchange fee. Clearly, the former mechanics lead to increased revenue for publishers as compared to the latter.

C. Last Look

157. "Last Look advantage" is a phrase used to refer to AdX's ability, under Dynamic Allocation and Enhanced Dynamic Allocation, to see the header bidding clearing price before submitting their own bid.[218] While the definition and mechanics of Dynamic Allocation do not change depending on whether other exchanges participated via the waterfall process or header bidding, the implications do. Here, I briefly elaborate on the related concepts.

158. I have previously described that Dynamic Allocation and Enhanced Dynamic Allocation offer a Last Look advantage to AdX when other exchanges participate in header bidding. AdX learns the highest header bidding bid before submitting its own. This is equivalent to a first-price auction where all bidders except for AdX submit their bids first, then AdX learns the highest submitted bid and submits its own.[219, 220]

1) Last Look helps AdX clear impressions that would have otherwise been cleared by the header bidding winner, by paying a penny more

159. Last Look advantage likely helped AdX have a higher win rate in comparison to AdX's win rate without Last Look advantage, by helping AdX clear impressions that would have otherwise been cleared by the header bidding winner.[221] I justify this claim below.

a. I consider two cases: (a) all exchanges submit bids via a first-price auction with personalized reserves, and (b) all non-Google exchanges submit bids via a first-price auction with the same personalized reserves, and AdX's Last Look

Server-side header bidding take rates vary by the service provider. For example, Ad Butler charges $100 per month and $0.001 per a thousand bids for its "Programmatic Advertising" service. AdButler. "Get the Industry-Leading Ad Server." Accessed on May 31, 2024.
https://web.archive.org/web/20231129061555/https:/www.adbutler.com/pricing.html
[218] And, for a period, the Exchange Bidding clearing price as well. GOOG-TEX-00000744 at -54. April 26, 2017. "Exchange Bidding (Jedi) Open Beta Sates Readiness Review."
[219] I have analyzed an example in Section IV.A.2 to demonstrate how the Last Look Advantage helps AdX in this situation.
[220] Even for sophisticated publishers, AdX learns some information about the highest submitted bid (and in particular, that the highest submitted bid is below its reserve price).
[221] This is my opinion in aggregate, after considering that some publishers used default options while others were sophisticated and increased Value CPMs of header bids to boost AdX's reserve.

advantage lets it win instead as long as it submits a bid exceeding both its (same) personalized reserve and the highest first-price bid.[222]

b. Let h denote the highest bid from a non-Google exchange that clears its reserve, v denote the highest value of any advertiser who bids with AdX, and r denote AdX's personalized reserve. There are a few cases to consider.

   i. If either h or r exceeds v, then certainly AdX will not win in either auction format. Therefore, both formats have the same outcome.

   ii. If v exceeds both h and r, then AdX certainly wins with Last Look advantage. This is because AdX's reserve will be the maximum of h and r, and its advertiser with value v will certainly submit a bid exceeding h and r,[223] and therefore some AdX bidder will win. If AdX instead participates with other exchanges in the first-price auction, AdX might lose. In particular, if AdX runs a second-price auction, their highest-value bidder will still bid v, but AdX would need a second highest bid (or reserve) exceeding h in order to win. Similarly, if AdX runs a first-price auction, their highest-value bidder would shade their bid, and without knowing h might shade their bid below h. Therefore, in any auction where AdX has a chance to win while participating in a first-price auction, AdX certainly wins with its Last Look advantage (and there are cases where AdX would win with Last Look but lose without it).

c. As a result, the Last Look advantage helps AdX win more impressions that would have gone to the header bidding winner otherwise, assuming publishers and

[222] Recall that I have shown in Appendix D a natural example where optimal bidders would submit the same bids in both settings, although this would not hold in all settings. Recall also that this justification covers a default publisher who sets the same personalized reserves and does not boost header bids. I will briefly discuss how this analysis changes for sophisticated publishers.

[223] This holds if AdX runs a second-price auction, because the highest-value advertiser will submit a bid of v. It also holds if AdX runs a first-price auction, because the highest-value advertiser will strictly prefer to submit a bid that has a chance of winning than a bid that will certainly lose (although it is possible that their bid will ultimately lose to a higher bid of another AdX bidder, some AdX bidder will certainly win).

advertisers who bid the same whether or not AdX has a Last Look advantage.[224, 225]

160. In addition, when AdX wins these additional impressions, in many cases it pays just a penny more than the winning header bid and so does not increase publisher revenue. Specifically, in cases where AdX's highest value v exceeds the highest header bid h, but h exceeds both AdX's second highest value and its reserve, then AdX wins the impression with Last Look at price h (plus a penny), while without Last Look the header bid of h would have won. When some exchanges participate in header bidding while others participate in exchange bidding,[226] only those who participate in header bidding are vulnerable to AdX's Last Look advantage.[227] Specifically, the highest header bidding bid would become the reserve for AdX, whereas bids through Exchange Bidding are not revealed to AdX before AdX submits its own bid. That is, within AdX and Exchange Bidding exchanges, no one has a Last Look Advantage over the other, because their bids are submitted simultaneously.[228] Exchanges that are integrated into Exchange Bidding see the same DFP reserve as AdX, hence they also have a Last Look advantage over header bidding exchanges.[229]

161. Therefore, one interpretation of Exchange Bidding is that it creates two tiers: Exchanges that participate in header bidding and exchanges that participate in Exchange Bidding together with AdX.[230] Exchanges that participate in header bidding submit bids without seeing others' bids, and therefore have no Last Look advantage over anyone, and are vulnerable to AdX's and Exchange Bidding's Last Look advantage (placing them in the lowest tier).[231] AdX and Exchange Bidding exchanges have a Last Look advantage over exchanges that participate in header bidding, but do not have a Last Look advantage over each other, placing them in the top tier.

[224] Recall again that I have shown a natural example in Appendix D where bidding identically in these situations is optimal for advertisers.
[225] If a sophisticated publisher mildly inflates AdX's reserve specifically because of the Last Look advantage, or if advertisers bid similarly in these two cases, the same conclusions still qualitatively hold. If a sophisticated publisher drastically inflates AdX's reserve specifically because of the Last Look advantage, or advertisers drastically change their bids specifically due to AdX's Last Look advantage, the impact is less clear-cut and would require a complicated analysis weighing the benefits of Last Look versus the impact of an increased reserve and distinct bids.
[226] Importantly, exchanges could participate in both. I am not claiming that this was the case, I am providing explanations based on a hypothetical.
[227] Last Look was phased out in 2019, when AdX transitioned to a first-price auction format. *See* GOOG-TEX-00841386 at -88. "Adx First Price Auction." ("removing last look.")
[228] *See* GOOG-TEX-00000744 at -54. April 26, 2017. "Exchange Bidding (Jed') Open Beta Sates Readiness Review." (diagram shows the '3p floor' entering AdX, but not other exchanges.)
[229] GOOG-DOJ-AT-01815211 at -222. October 2019. "Open Bidding (fka Exchange Bidding)Training."
[230] Again, exchanges can happen to be in multiple of these groups, since they can integrate into both header bidding and Exchange Bidding.
[231] As previously noted, if exchanges that are integrated into header bidding see the same DFP reserve as AdX, header bidders are also vulnerable to a Last Look from these exchanges.

Because a Last Look advantage is significant in first-price auctions, it would be natural for exchanges to want to remove the top tier's Last Look advantage over them and to gain a Last Look advantage over header bidders,[232] even though DFP takes a 5% fee on top of the clearing price when the winner is an Exchange Bidding exchange.[233, 234]

VI. CONDUCT ANALYSIS: UNIFIED PRICING RULES

162. In this section, I provide an analysis of Google's Unified Pricing Rules (UPR), which was instated in 2019 (and in place today[235]) along with Google's ad exchange AdX's transition to the first-price auction format.[236]

163. I demonstrate that UPR leads to lower revenue for the publishers. I also demonstrate that UPR can lead to better win rate and revenue for Google's ad exchange AdX as well as for Google's ad buying tools and lower the win rate and revenue for rival exchanges and ad buying tools.

164. Prior to UPR, publishers could set different reserves that applied to different exchanges or different ad buying tools. Under UPR, publishers can no longer employ these personalized reserves to their full extent,[237] because any reserve price set for non-guaranteed line items applies to all non-guaranteed line items.[238] This reduces publisher choice by preventing them from setting personalized reserve prices. Publishers retain the ability to set personalized reserves on individual advertisers, but not on individual exchanges or ad buying tools.[239] Furthermore, publishers may

---

[232] And after the change referenced in GOOG-DOJ-AT-01809483 went live, it would be further natural for exchanges to want a Last Look over header bidders. GOOG-DOJ-AT-01809483 at -89. March 2017. "Exchange Bidding in Dynamic Allocation (fka Project Jedi)."
[233] Of course, the most natural auction format is to avoid creating tiers and a Last Look advantage at all, and to simply have all exchanges submit bids simultaneously without seeing each other's.
[234] Last Look advantage was removed in 2019 during the implementation of Unified Pricing Rules and AdX's switch to the first-price auction format. *See* GOOG-TEX-00841386 at -89. "Adx First Price Auction."
[235] Google. "Unified pricing rules." Accessed on May 31, 2024. https://web.archive.org/web/20230208153751/https:/support.google.com/admanager/answer/9298008?hl=en (current Google Ad Manager documentation on UPR).
[236] *See generally* Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20240122142404/https:/blog.google/products/admanager/update-first-price-auctions-google-ad-manager/
[237] GOOG-AT-MDL-000875073 at -83. August 2019. "The Unified First Price Auction."
[238] *See* Google. "Unified pricing rules." Accessed on May 31, 2024. https://web.archive.org/web/20230208153751/https:/support.google.com/admanager/answer/9298008?hl=en (current Google Ad Manager documentation on UPR).
[239] More specifically, publishers are free to set their reserve price at any level they desire, however this reserve price applies to all exchanges and ad buying tools.

set at most 200 total reserve prices at the advertiser level.[240, 241] The excerpt in Figure 30 from an internal Google slide deck[242] discusses the reserve setting abilities publishers lose under UPR.

Figure 30: An excerpt from an internal Google document specifying that reserve prices under UPR applies to all non-guaranteed line items[243]

Summary of changes with Unified Pricing Rules

	AdX Open Auction Pricing Rule	Unified Pricing Rule
Floor applies to	• Authorized Buyers (Ad Exchange)	• Authorized Buyers (Ad Exchange) • Exchanges on EBDA • Non-guaranteed line items (excluding $0 non-guaranteed and House)
Rules priority, overlapping rules	Floor from rule with **higher priority** will apply to the demand available	**Maximum available floor** will apply to the demand available
Branding Types	Branded, Semi-Transparent, Anonymous	Branded, Semi-Transparent (no different pricing per branding type)
Per- buyer floor	Yes	No
Blocks (buyer / advertiser)	Pricing Rule UI	Blocks migrate to *Protections* UI *

* Coming soon

Google

165. As I outlined in Section II.B, there are several possible reasons why publishers might choose to set personalized reserve prices. First, they may wish to compensate for ad quality. A publisher may wish to charge more for low-quality ads (whose display negatively impacts user experience and indirectly causes the publisher financial loss) than high-quality ads (whose display

[240] More precisely, any UPR reserve applies to all exchanges in Exchange Bidding as well as AdX. They also apply to line items that correspond to the header bidding winning bids. Note that publishers can configure any floors they like within their header bidding setup, including personalized reserves, but the header bidding line item will still be blocked if they are under the relevant UPR reserves. *See generally* Google. "Unified pricing rules." Accessed on May 31, 2024. https://web.archive.org/web/20230208153751/https:/support.google.com/admanager/answer/9298008?hl=en (current Google Ad Manager documentation on UPR).
[241] At the initial announcement stage, Google stated that the system would allow 100 price floor rules. *See* Sarah Sluis. "Publishers Lash Out Against Google Over 'Unified Pricing' Changes" (April 18, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20221102060335/https:/www.adexchanger.com/online-advertising/publishers-lash-out-against-google-over-unified-pricing-changes/; News Corp Australia. "Submission To the Australian Competition and Consumer Commission" (May 2020). Accessed on May 31, 2024. https://web.archive.org/web/20221012074940/http:/www.accc.gov.au/system/files/News%20Corp%20Australia%20(15%20May%202020).pdf But based on the negative feedback, they increased the rule limit to 200. *See* GOOG-TEX-00594205 at -11. "The Unified First Price Auction Best Practices." ("Based on your feedback, we are improving Unified Pricing Rules; Coming soon; Unified pricing rules limit increased to 200.")
[242] GOOG-AT-MDL-000875073. August 2019. "The Unified First Price Auction."
[243] GOOG-AT-MDL-000875073 at -83. August 2019. "The Unified First Price Auction."

might be neutral, or at least less negative, to user experience). If a publisher believes that ads served through AdX on average cause more financial loss in comparison to ads served through other exchanges, it makes sense to set a higher personalized reserve to AdX than to other exchanges. Furthermore, publishers can block ads coming from an exchange altogether if they deem the ads particularly damaging, by setting a very high reserve for that exchange. Personalized reserve prices allow the publisher to express their financial preference for different ads and improve efficiency. Google is also aware of the publishers' attention to ad quality, stating that "many AdX publishers are very sensitive to ads that they feel reflect badly on their brand and unwilling to risk any exposure at all; when such ads […] appear on their pages, they may react by pulling their inventory from AdX altogether."[244]

166. Another reason why publishers may wish to set personalized reserves is that they might want to charge more to an exchange that tends to produce higher bids since these exchanges will generate higher revenue for the publishers. One exchange might tend to produce higher bids than another for two key reasons: (a) the exchange simply has access to a larger number of advertisers[245] or (b) the exchange has behavior that causes them to submit higher bids with access to a comparable advertiser number (such as Google's Project Bernanke, which I analyze in Section VIII).

1) UPR can lead to inefficient outcomes

167. Imagine a publisher finds that ads served through AdX to on average cause financial loss of $5 more than ads through OpenX do, and therefore chooses to set a personalized reserve $5 higher on AdX (to have concrete numbers in mind, say that OpenX causes $2 loss on average, while AdX causes $7 on average).[246] Hence the publisher sets a reserve of $15 for AdX and $10 for OpenX. AdX submits a bid of $13, and OpenX submits a bid of $11. Then with personalized reserves, OpenX will win the impression, and this is the efficient outcome (yielding payoff $9 for the publisher; $11 of revenue minus $2 of harm). Under UPR, the publisher instead must set the same reserve on both exchanges. That reserve might be less than or equal to $13 which will result in AdX winning the impression, which nets the publisher more revenue. But the publisher's

[244] GOOG-DOJ-15769995 at -5. May 2017. "Protecting Publishers from Objectionable Ads - Proposal."
[245] And this itself might be due to 'good business' of providing a better service, or due to some of the other conduct allegations in this case where Google products preference AdX on both the buy-side (causing advertisers to end up using AdX even though another exchange might provide a better price or product) and the sell-side (causing the same bids through AdX to have a better shot at winning) of the market.
[246] Note that setting optimal reserves is a complex problem and setting a reserve that is exactly $5 higher on AdX because ads cause exactly $5 more harm on average is not likely to be optimal. However, it is a sensible heuristic, and the key takeaways from the example are not driven by the particular choice of personalized reserve.

true gain is $6; $13 of revenue minus $7 of harm. The publisher might instead set a single reserve that is higher than $13 which will result in no sale. In this case, there is no single reserve price which results in the efficient outcome of OpenX winning the impression, because UPR removes the ability of the publisher to set personalized reserve prices. This example is illustrated in Figure 31 below.

Figure 31: The efficient outcome cannot be achieved in this auction due to UPR

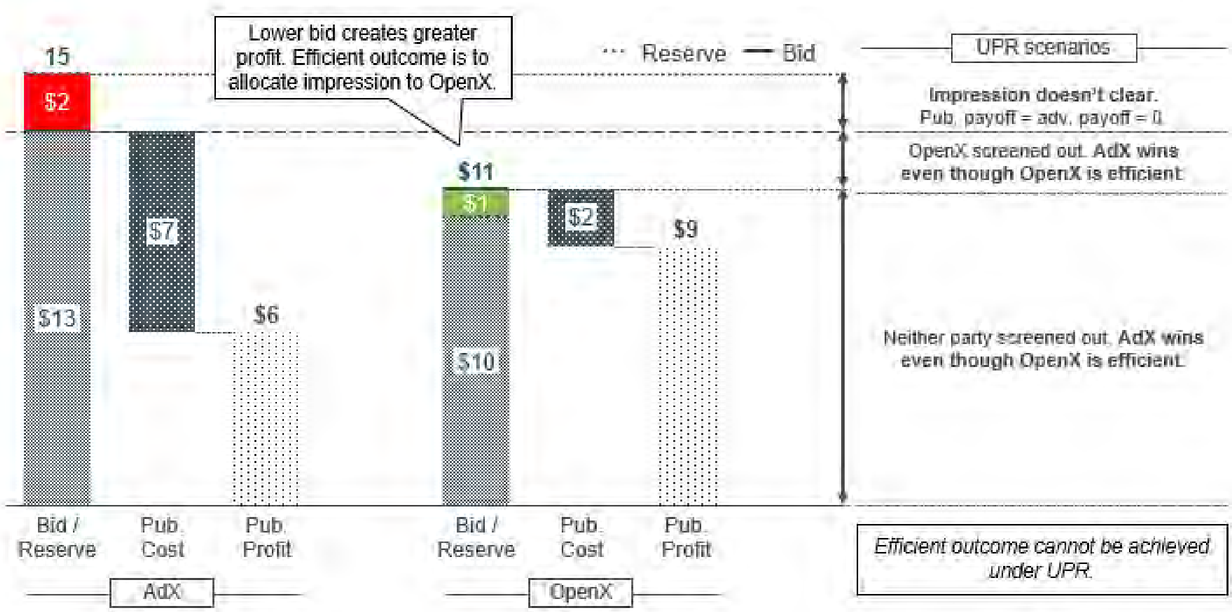

168. Optimal reserves improve revenue by putting pressure on bidders to increase their bids.[247] Imagine that the publisher knows that the top bid from AdX tends to be between $12 and $16, whereas the top bid from OpenX tends to be between $10 and $12, and for simplicity of this example assume that each exchange has a single advertiser.[248] The publisher might therefore set a reserve of $13 on AdX and $10 on OpenX. Under these reserves, the AdX advertiser would optimally shade any bids above $13 down to $13, and the OpenX bidder would optimally shade all bids down to $10. When AdX's bidder has a value of $15, these personalized reserves and

[247] Because UPR was implemented as AdX switched to a first-price auction format, and during a time when the majority of the intermediated display advertising ecosystem was also using first-price auctions. *See* Sarah Sluis. "Google Switches To First-Price Auction" (March 6, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20220910040643/https:/www.adexchanger.com/online-advertising/google-switches-to-first-price-auction/ ("Google Ad Manager will be the last major exchange to switch to first-price auctions. Other exchanges tested or rolled out first-price auctions starting in 2017.")
I focus my analysis on this case.

[248] To make this example mathematically rigorous, imagine that there is a single AdX bidder whose value is distributed uniformly on [12,16], and there is a single OpenX bidder whose value is distributed uniformly on [10,12]. In this case, the revenue-optimal personalized reserves are $13 on AdX and $10 on OpenX. Moreover, the example bids for AdX and OpenX are optimal in this case.

optimal advertiser behavior results in a revenue of $13. When AdX's advertiser instead has a value of $12, these personalized reserves and optimal advertiser behavior results in a revenue of $10, even in cases where AdX's advertiser does not clear the reserve price. But under UPR the publisher must set the same reserve for both exchanges. If that reserve exceeds $12, then it likely excludes OpenX,[249] and acts as a reserve on AdX alone. In particular, if AdX fails to meet the reserve, the publisher likely foregoes the fallback option of $10 they were able to achieve with personalized reserves. If instead that reserve is below $12, it puts no pressure on AdX's advertiser, which would optimally shade its bid down to $12.

169. In summary, personalized reserves allow the publisher to both set a competitive reserve for higher-priced exchanges and still collect some revenues when higher-priced exchanges fail to meet those reserves, while UPR forces the publisher to choose between one or the other. Furthermore, personalized reserves allow the publishers to screen for ad quality, by either charging more to compensate for the low quality or blocking the low quality ads altogether. UPR disables publishers from effectively using reserve prices to screen for ad quality.

B. Impact of Unified Pricing Rules on Publishers

170. In this subsection, I explain why UPR leads to a revenue loss for publishers. Under UPR, publishers lose the ability to maximize their revenues because they cannot set personalized reserves for exchanges and ad buying tools.

1) UPR prevents publishers from maximizing their revenues

171. The ability to set personalized reserves for each exchange is an important revenue-optimization tool for publishers. UPR took away this ability, as discussed above. This is true when publishers are interacting with an advertiser for the first time, as well as with an advertiser they have sufficient information about.

a. For new advertisers, or advertisers with whom the publisher has not interacted enough to be able to set a suitable reserve price, publishers would rely on coarse information to set appropriate reserves. For example, advertisers that choose to transact through AdX may tend to be materially different (both in willingness to pay and in ad quality) than advertisers who chose to transact through non-Google exchanges. In order to maximize revenue, publishers would still want to set as

[249] And in the mathematically rigorous example, completely excludes OpenX.

good a reserve as possible on new advertisers for whom they have not yet decided on an appropriate advertiser-specific reserve. The natural mechanism by which to do this is to set a single personalized reserve on AdX (or the non-Google exchanges) that applies to all advertisers, including new ones, who transact through AdX. Note that the ability to set personalized reserves to advertisers does not address this source of revenue loss since per-advertiser reserves are only meaningful with sufficient information about the advertiser, whereas a per-exchange reserves are necessary to maximize revenues based on available information on new advertisers.

b. Even when the advertiser is known to the publisher, publisher revenue can depend on the exchange through which that advertiser's bid occurred. The same advertiser's initial bid is processed differently through AdX than non-Google exchanges.[250] Therefore, a revenue-maximizing publisher would naturally want to set different reserves for exchanges that engage in different conducts, even if they reach the same advertiser pool.

c. Finally, even when the advertiser is known to the publisher, *and* that advertiser is known to transact through exactly one exchange, the '200 Rules' limit that comes with UPR would still lead to revenue loss among publishers because it prevents the implementation of revenue-maximizing strategies that use personalized reserves. When there are many advertisers potentially interested in an impression, revenue-maximization requires a granular reserve for each advertiser, as explained above. Limiting the number of allowed reserve prices to 200 limits the publisher's ability to maximize revenues.

2) UPR can lead to lower quality ads for the publishers

172. As discussed above, UPR can decrease the ad quality that the publishers face. This is because publishers are not able to filter for ad quality across different exchanges by employing personalized reserves. [redacted]

[250] For example, GDN's Project Bernanke inflates the highest bid sent to AdX. Therefore, the same advertiser's bid is likely higher after being processed through AdX than a non-Google exchange.

[redacted][251]

173. In sum, UPR would lead to a revenue loss for publishers. UPR would decrease publisher welfare even further by preventing them from effectively screening for low quality ads.

C. Impact of Unified Pricing Rules on Exchanges and Ad Buying Tools

174. In this subsection, I demonstrate that UPR can lead to a better win rate and higher revenue for Google's ad exchange AdX as well as Google's ad buying tools, and lower win rate and revenue for rival exchanges and ad buying tools. I provide explanations as to why these are plausible outcomes of UPR. I show that if AdX was facing high reserves compared to other exchanges compared, then it is plausible to expect that AdX revenue increased as a result of UPR. Furthermore, I explain that UPR can hurt exchanges and ad buying tools that facilitate the transaction of high quality ads.

175. Under UPR, for each advertiser, AdX receives a reserve price that is no higher than any other non-guaranteed line item, including other exchanges. This is true for exchanges that are called via Exchange Bidding, the waterfall[252] or header bidding.[253] In other words, for a given impression, all exchanges face the same reserve. This can benefit AdX if it previously faced higher reserves compared to other exchanges.

176. Since setting personalized reserves is a complex problem, it is likely that publishers use general heuristics.[254] One plausible such heuristic is the **Treat-As-Single Heuristic**, where publishers set personalized reserves for each individual exchange as if it were the only exchange (optimally according to this assumption, using Myerson's (1981)[255] reserve),[256] and publishers set

[251] GOOG-AT-MDL-004017152. ("[redacted].)
[252] Note that exchanges called through Exchange Bidding must pay the Exchange Bidding fee, and that exchanges that participate through the waterfall may not be called at all.
[253] Exchanges that are called through header bidding face a reserve price at least as high as the UPR reserves but may face a higher personalized reserve. This is because UPR applies to all non-guaranteed line items, which include header bidding line items, but the header bidding technology still allows for personalized reserves, so publishers can still choose to set personalized reserves in the header bidding auction, with the condition that these must be higher than the applicable UPR reserve set in DFP.
[254] A heuristic for a publisher in this case can be thought of a guideline that determines their actions.
[255] Roger B. Myerson. "Optimal Auction Design." *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73. Subsequent works are also relevant.
[256] The personalized reserves set under the Treat-As-Single Heuristic are a commonly studied heuristic in auction theory under the name "VCG with Monopoly Reserve." *See, e.g.*, Hartline and Roughgarden. "Simple versus Optimal Mechanisms." *Proceedings of the 10th ACM Conference on Electronic Commerce*. 2009. pg. 225-234.

UPR reserves by setting Myerson's optimal reserve for a hypothetical "average exchange."[257] Another plausible heuristic is the **eCPM Heuristic**, where publishers set personalized reserves for each individual exchange equal to the exchange's historical average bid on similar impressions (after adjusting for value), and publishers set UPR reserves equal to the average historical bid on similar impressions (across all exchanges, after adjusting for value).[258] [259]

177. If publishers use the Treat-As-Single Heuristic, then whatever exchange faces the highest personalized reserve would face a lower reserve under UPR. I provide a proof of this claim in Appendix F. The proof establishes that the Treat-As-Single Heuristic sets a UPR reserve that is somewhere between the minimum and maximum personalized reserves, under a standard 'regularity assumption' on the bids typically produced by each individual exchange. Therefore, whichever exchange faces the highest reserve will also face a lower UPR reserve, and whichever exchange faces the lowest reserve will also face a higher UPR reserve. If publishers use the eCPM heuristic, then whatever exchange faces the highest personalized reserve would also face a lower reserve under UPR. The argument again establishes that the eCPM heuristic sets a UPR reserve somewhere between the minimum and maximum personalized reserves (and does not require a regularity assumption). Therefore, again, whichever exchange faces the highest personalized reserve will also face a lower UPR reserve, and whichever exchange faces the lowest reserve will face a higher UPR reserve. While it is too complex to guess precisely how publishers set reserves, these two natural heuristics demonstrate similar behavior and result in UPR reserves somewhere between the minimum and maximum personalized reserves, which implies lower UPR reserves for exchanges that received the highest personalized reserves. It is therefore natural to expect other publisher heuristics to follow a similar analysis.

1) If AdX faced the highest personalized reserve pre-UPR, then AdX would transact more impressions and have increased revenue under UPR

178. If AdX faced the highest personalized reserve pre-UPR, then AdX would likely transact more impressions and have increased revenue under UPR. There are two steps to understand why this is the case. First, if AdX faced the highest personalized reserve pre-UPR, then it is natural to expect AdX's UPR reserve to be lower, as I explained previously. Second, I now argue that

---

[257] To be clear, by "average exchange," I mean that the publisher considers an exchange with value distribution whose cumulative distribution function is the average of the cumulative distribution functions of all individual exchanges. In more formal terms, if exchange *i* has cumulative distribution function F_i, then the "average exchange" of *n* exchanges has the cumulative distribution function $\sum_{i=1}^{n} F_i/n$.
[258] To be clear, by "adjusting for value" I mean the following: "If the harm caused by displaying ads from this exchange is on average $X CPM, set the reserve for this exchange equal to its eCPM plus $X."
[259] I give more detailed definitions for these heuristics in Appendix F.

AdX would transact more impressions and have increased revenues under UPR, assuming that indeed AdX's UPR reserve is lower than its pre-UPR personalized reserves. Specifically, I now compare the win-rate and competition faced by a particular exchange in the following two scenarios: (a) Personalized reserves, with this particular exchange's being the highest, and (b) unified reserves, and less than this particular exchange's personalized reserve in (a). I conclude that, if the bidder behavior results in efficient outcomes,[260] this particular exchange:

a. Wins more impressions under unified reserves. More specifically, for each impression that this particular exchange would have won without UPR, this particular exchange also wins under UPR. In addition, there are impressions this particular exchange would not have won without UPR that this particular exchange now wins under UPR.

b. Faces less competition faced for impressions it wins on average under unified reserves. That is, for the set of impressions that this particular exchange would have won without UPR, this particular exchange also wins these impressions under UPR. Moreover, the competition faced on these impressions is no higher under UPR than without UPR.[261]

I elaborate on these points further in Appendix F.

179. Any exchange whose pre-UPR personalized reserve is lower than their UPR reserve would experience a lower yield and lower revenues under UPR, in comparison to the counterfactual of not imposing UPR on publishers. Moreover, if AdX faced the highest personalized reserve pre-UPR, some other exchanges must face a higher UPR reserve than their pre-UPR personalized reserves under either of the two natural heuristics, and therefore would achieve a lower yield and revenues under UPR in comparison to the counterfactual of not imposing UPR on publishers. I compare the outcomes for exchange *i* under two sets of reserves: (a) Personalized, where exchange *i* faces a reserve of r_i, and (b) a unified reserve of $r \geq r_i$. I conclude that, if the bidder behavior results in efficient outcomes, Bidder *i* would:

[260] The outcome of an auction is efficient if and only if, among the bidders who submitted a bid higher than their personalized reserves, the bidder with the highest valuation of the auctioned item wins the item.
[261] Formally, what I mean by this is that treating the exchange as a single bidder, that exchange *could* pay less on average and still win the same impressions under UPR as pre-UPR. The exchange may still choose to pay the same or more (i.e., by setting a reserve to its advertisers that exceeds the competition faced, using a program like Google's Reserve Price Optimization). Or the exchange may instead choose to charge its advertisers the same as when it would win pre-UPR, and simply pass on less payout to the publisher (i.e., using a program like Google's Dynamic Revenue Sharing, although that family of conducts never occurred at the same time as UPR).

a. Win fewer impressions under UPR. That is, for each impression that exchange *i* would have won under UPR reserve of *r*, exchange *i* certainly wins under personalized reserves of r_i.

b. Faces higher competition for the impressions they win on average. That is, for the set of impressions that exchange *i* wins under UPR at reserve *r*, exchange *i* also wins these impressions at personalized reserves r_i. Moreover, the competition faced on these impressions is no higher under personalized reserves r_i than under UPR at reserve *r*.[262]

I elaborate further on these points in Appendix F.

2) UPR would decrease the win rates and revenues of exchanges and ad buying tools that typically transact high quality ads

180. UPR would decrease the win rates and revenues of exchanges and ad buying tools that typically transact high quality ads. In general, UPR requires that bids are treated identically from all ad buying tools and all exchanges. In contrast, personalized reserves allow publishers to preference bids of the same dollar value from ad buying tools and exchanges that transact higher quality ads. Under personalized reserves, it is plausible that publishers set higher personalized reserves for ad buying tools and exchanges that facilitated transactions of lower quality ads in comparison to others, which might enable, for example, a bid of $5 from a high-quality ad to be selected over a bid of $5.01 from a low-quality ad. But, under UPR, all bids are treated identically and a $5.01 bid from a low-quality ad will always be selected over a $5 bid from a high-quality ad.[263] Therefore, exchanges that typically transact high-quality ads would have lower win rates and revenues under UPR.[264]

181. For example, if the typical values of bidders are identical across exchanges, but the typical ad quality differs, the exchange with the highest ad quality would face a lower reserve with personalized reserves compared to UPR, under natural heuristics such as the Treat-As-Single

[262] Again, what I formally mean by this is that treating the exchange as a single bidder, that exchange *must* pay more on average to win any impression that it wins under UPR as compared to what it must pay pre-UPR. As noted previously, the exchange can always choose to set a higher reserve on its advertisers than it must pay, so all exchanges prefer to be less constrained in what they *must* pay to a publisher in order to win an impression.
[263] The lone exception to this is if the $5.01 advertiser is subject to one of the 200 rules. As previously noted, 200 rules may be sufficient to screen out some undesirable advertisers but is insufficient to express a general preference for high-quality over low-quality ads in the same manner as personalized reserves.
[264] Due to the complexity in optimizing reserves and equilibria of first-price auctions, it is challenging to make strong predictions about what might happen in practice from an auction theory perspective alone. The analysis states what I would see as the most likely outcome in such a scenario, based on my expertise in auction theory.

Heuristic or the eCPM heuristic. Exchanges with high ad quality would therefore suffer lower win rates and revenues under UPR. Between two exchanges with identical typical values but distinct typical ad quality, the exchange with higher ad quality would face a lower Myerson reserve when treated as the only exchange, and this would result in higher win rates and higher revenues under the Treat-As-Single Heuristic. Similarly, between two exchanges with identical typical values but distinct typical ad quality, the higher ad quality would have a stronger value-adjusted eCPM (because the eCPMs are identical, but value-adjusting works more strongly in favor for higher-quality ads). I provide a more formal explanation for this result in Appendix F.

182. In sum, assuming that AdX faced the highest reserve pre-UPR, UPR would naturally benefit Google's ad exchange AdX both in win rate and in revenue.[265] Furthermore, it negatively affects some non-Google exchanges and ad buying tools, especially those that tend to transact high quality ads.

VII. CONDUCT ANALYSIS: DYNAMIC REVENUE SHARING

183. In this section, I analyze Google's Dynamic Revenue Sharing (DRS)[266] conduct and its variants, which went on from 2015 to 2019. I find that:

a. Dynamic Revenue Sharing version 1 (DRSv1) increased AdX win rate and revenue and decreased non-AdX exchanges' win rates and revenues, compared to no DRS,

b. Dynamic Revenue Sharing version 2 (DRSv2), in comparison to both no DRS and DRSv1, decreased advertiser payoff, increased AdX win rate and revenue, decreased non-AdX exchange's win rates and revenues, and may also decrease publisher revenue.

c. Google concealed information that is vital to advertisers and important to publishers by concealing DRSv1 from them. At least some of Google's communication regarding DRSv2 was misleading.

[265] If AdX faced a relatively high reserve pre-UPR, but not the *highest* reserve pre-UPR, all of my conclusions qualitatively hold for the same reasons. I chose to present the case where AdX faces the *highest* reserve in the interest of crisp statements and clean proofs.
[266] This conduct is sometimes referred to as "sell-side DRS." GOOG-AT-MDL-003849201. Email thread, "Subject: Re: Partner revshare scaling factor alerts." ("The timing seems to match with the sellside DRS experiment.")

184. Before the DRS program, Google imposed a contractually mandated[267] per-auction take rate that was typically equal to the 20% of what is submitted by the ad buying tool.[268] This implies that Google either (a) took 20% of the second highest net bid[269] and passed on the remaining 80% to the publisher if the remaining 80% of the second highest net bid was higher than the publisher's reserve or (b) paid the publisher their reserve and charged the advertiser an amount that corresponds to 125% of the reserve,[270] since the 80% of that amount would exactly be the reserve.

185. As a result, prior to DRS, AdX ran the following auction:

a. AdX learns its reserve r from the ad server. Due to the static take rate of 20% imposed by AdX, the highest bid submitted to the AdX auction must be at least $r/0.8$ to win the impression.[271]

b. AdX solicits bids from the ad buying tools. Let b_1 denote the highest received bid and b_2 denote the second highest.[272]

c. If the second highest bid is above the effective reserve[273] (*i.e.*, $b_2 \geq r/0.8$), then the clearing price is b_2, AdX charges $0.2*b_2$ as its fee, and passes on $0.8*b_2$ to the publisher.

d. If the effective reserve is in between the highest and second highest bids (*i.e.*, $b_1 \geq r/0.8 > b_2$), then the clearing price is $r/0.8$, AdX charges $0.2*r/0.8 = r/4$ as its fee and passes on $0.8*r/0.8 = r$ to the publisher.

e. If the highest bid is below the effective reserve (*i.e.*, $r/0.8 > b_1$), then the highest bid fails to clear the reserve price after AdX takes its rate, hence AdX fails to win the impression.

---

267 [redacted].")
268 [redacted].")
269 That is, the bid submitted by the ad buying tool, after having already taken the ad buying tool fee.
270 Abstracting away from the ad buying tool fee.
271 AdX's take rate is taken based on the clearing price.
272 For the sake of clarity, I abstract away from the ad buying tool fee, since it is immaterial to the analysis conducted here.
273 By "effective reserve price," I mean the reserve price plus the take rate AdX take if the reserve price ends up being the clearing price.

186. This is a regular second-price auction in which bids are adjusted by the exchange's take rate. Until 2019, Google stated that it ran a second-price auction.[274] The second price auction is truthful, which means that it is in the advertiser's best interest to bid their true value for the impression. This has many desirable properties, as noted in Section II. I note that the first price auction is not truthful.[275]

187. In order to analyze the impacts of the different types of DRS conducts, I first introduce a **dirty second-price auction**, as defined in Google's internal documentation.[276] A dirty second-price auction has two parameters, a hard floor r, and a soft floor $q \geq r$. There are many ways to implement a dirty second-price auction, but they will all have the following properties when only the highest bid b_1 exceeds the hard floor r.

a. If the highest bid is above the soft floor (*i.e.*, $b_1 \geq q$), then the highest bidder wins and pays q. This is the same as the regular second price auction with reserve q.

b. If the highest bid is below the hard floor (*i.e.*, $b_1 < r$), then no one wins. This is same as the regular second price auction with reserve q.

c. If the highest bid is in between the soft and hard floors (*i.e.*, $q > b_1 \geq r$), then the highest bidder wins and pays b_1. This differs from the regular second price auction with any reserve, due to the fact that the highest bidder is paying their own bid. Hence for this specific case, the dirty second price auction gives the same outcome as a first price auction.

188. Notice that with dirty second price auctions, under some conditions, the highest bidder pays their own bid. This is the same as the first price auction. Since the first price auction is not truthful, this means that the dirty second price auction is not truthful.[277] To see this, suppose there is a bidder who has a value v that exceeds r but not q, and that no other bidder has a value exceeding r. If this bidder bids their true value, they will win the impression and pay v. If instead

[274] *See* Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20240122142404/https:/blog.google/products/admanager/update-first-price-auctions-google-ad-manager/
[275] See Section II for more information about truthful auctions.
[276] GOOG-NE-05279363 at -73. "Bidding in adversarial auctions."
[277] Under a dirty second price auction, the bidders have an incentive not to be truthful only under some conditions, but a truthful auction format is defined as always incentivizing the bidders to bid their true values. Since the dirty second price auction fails to do that under some conditions, it is fails to be truthful in general.

they submit a bid, b, between the reserve, *r*, and their value, v, they will win the impression and pay b < v which is strictly better than bidding their true value.

A. Dynamic Revenue Sharing v1

189. DRS was launched in August 2015[278] without announcing it to publishers or advertisers.[279] In its first iteration, DRSv1, AdX dynamically decreased its take rate to be lower than 20% to win impressions that it would not have if the take rate was kept at 20%. AdX decided to either impose a 20% take rate or decrease the take rate depending on a few factors including comparisons between the first and the second highest bid and the publisher reserve, [280] as well as the average take rate among the auctions for that publisher's impressions in that billing period.[281]

1[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[278] GOOG-TEX-00777528 at -30. Email thread, "Subject: Re: [Monetization-pm] Re: [drx-pm] LAUNCHED! AdX Dynamic Revenue Share (DRS)." (The email from September 2nd, 2015 states that "Last week we launched Dynamic sell-side Revenue Share (DRS).")

[279] [redacted]

[280] GOOG-NE-06864639 at -43. May 9, 2014. "Dynamic Sell-side Revshare on AdX." (under the subsection "Cases.")

[281] [redacted]

[282] This conduct occurred exclusively during a period where AdX ran a second-price auction.

[283] For the sake of clarity, I abstract away from the ad buying tool fee, since it is immaterial to the analysis conducted here.

i. [redacted]

[redacted]

[redacted]

191. DRSv1 impacts the auction only in one case where the highest bid is high enough to clear the reserve price, but not high enough to do so if AdX takes its full fee of 20% of the clearing price. When that happens, AdX dynamically decides to decrease its take rate so that it returns a successful bid to the ad server and wins the impression.

192. To illustrate how DRSv1 works, imagine an impression arrives, and the publisher reports a price floor of $10.[288] AdX solicits bids and receives top-two bids of $20 and $10. In this case, with both a regular second-price auction and DRSv1, the AdX clearing price is $12.5, because $20 ≥ $12.5 > $10. AdX takes a 20% take rate of $2.5 and passes on $10 to the publisher and DRSv1 has no impact. This example is illustrated in Figure 32 below.

[redacted]

288 For ease of exposition, I use ‘price floor’ for when the publisher sets a reserve on an exchange, and ‘reserve’ for when an exchange sets a reserve on advertisers.

Figure 32: An auction where DRSv1 has no impact

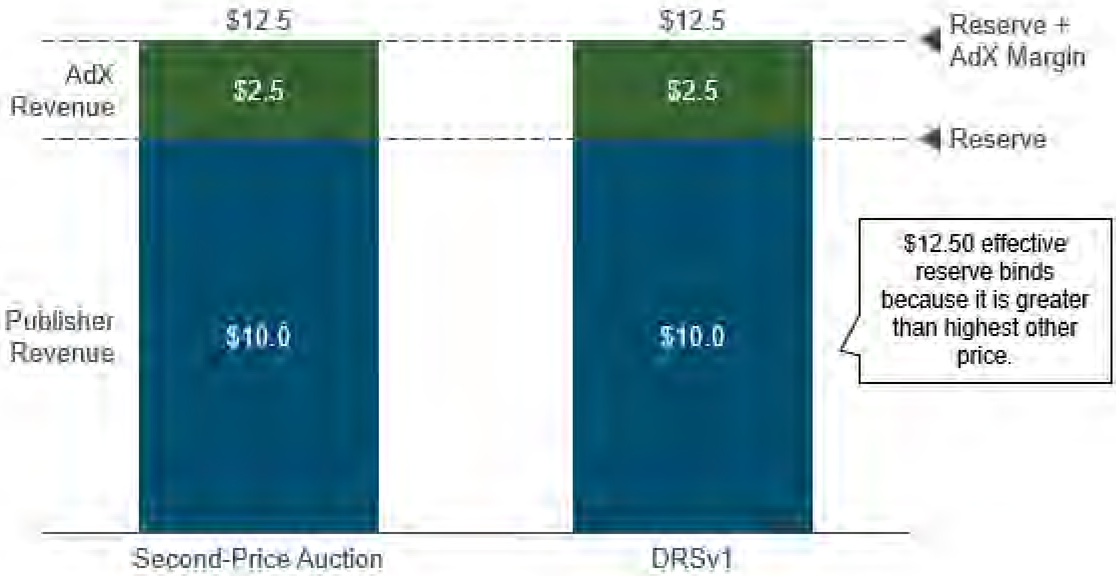

193. Imagine another impression arrives, and the publisher again reports a price floor of $10. AdX solicits bids and this time receives top-two bids of $9 and $8. In this case, with both a regular second-price auction and DRSv1, the impression is not transacted with AdX, because even with a take rate of 0% AdX still cannot clear the reserve.[289] This example is illustrated in Figure 33 below.

[289] This is a key distinction between DRS and Project Bernanke. Project Bernanke in some sense is a 'step further' than DRS because it further tries to transact at a negative margin. See Section VIII for further details on Project Bernanke.

Figure 33: An impression that AdX could not clear even with DRSv1

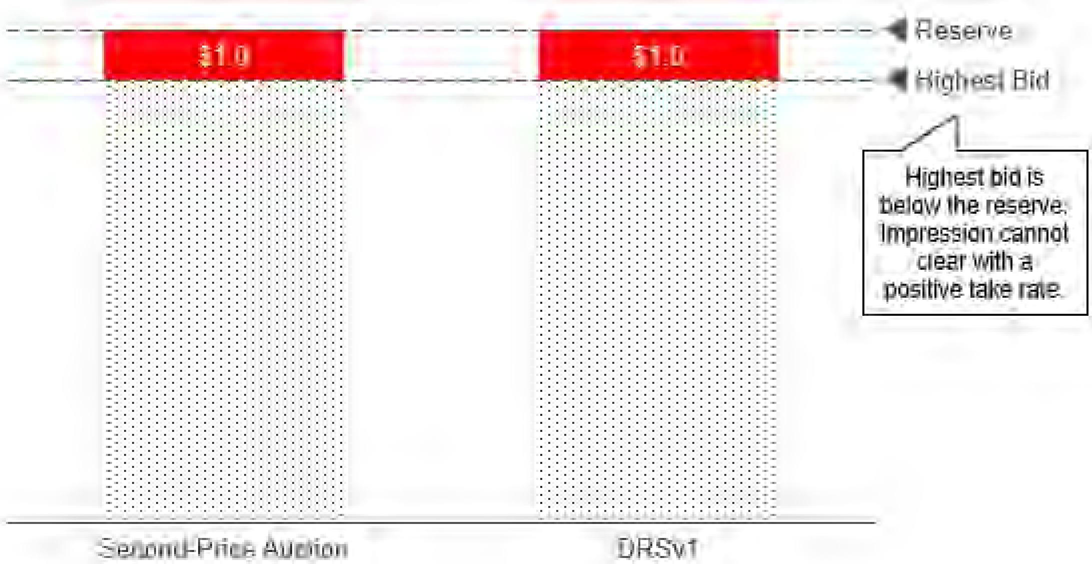

194. Imagine a final impression arrives, and the publisher again reports a price floor of $10. AdX solicits bids and this time receives top-two bids of $11 and $10. Under a regular second-price auction with a take rate of 20%, the impression would not be transacted with AdX, because that would require a top bid of at least $12.5 for a 20% take rate to still result in a payment above $10. Under DRSv1, provided that DRSv1 is not throttled,[290] DRSv1 would instead clear the transaction, charge the winner $11, collect a $1 fee for a take-rate of 9.1%. and pass on $10 to the publisher. Note that this lowers the average take-rate across the billing period.[291] This example is illustrated in Figure 34 below.

[290] If DRSv1 is throttled, DRSv1 will just execute as a regular second-price auction.

[291] In the final example, it is important to note that the advertiser is paying their bid. In particular, if the advertiser had reported a bid of $12 instead, they would have been charged $12 (and AdX would have taken $2, for a take-rate of approximately 16.66%). If the advertiser had reported a bid of $10.1, they would have been charged $10.1 (and AdX would have taken $0.1 CPM, for a take-rate of approximately 1%). That is, while DRSv1 acts like a truthful second-price auction for some bids, it is instead a first-price auction in the 'dynamic range' from $10 to $12.5, where AdX will charge the winner their bid.

Figure 34: An impression cleared by AdX with the DRSv1 dynamic adjustment of the take rate

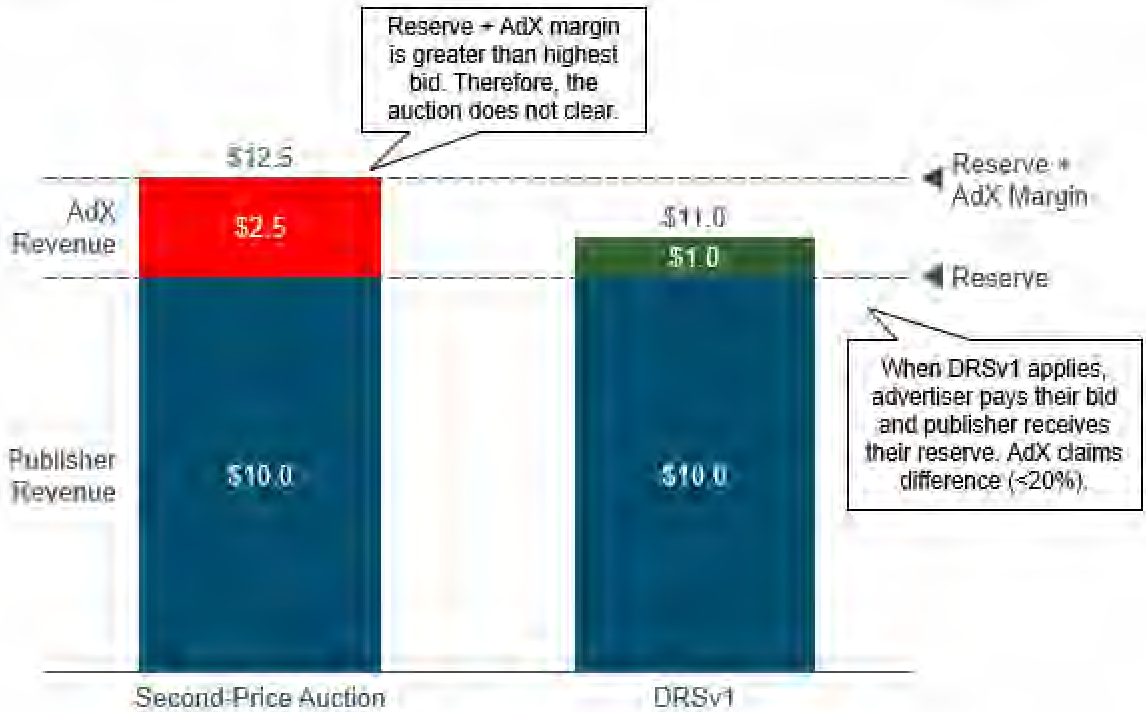

195. DRSv1 is a dirty second-price auction and as a result, DRSv1 is not truthful. Specifically, whenever DRSv1 charges the highest bidder their bid and pays the reserve r to the publisher, it is a dirty second-price auction with a hard reserve of *r* and a soft reserve equal to *r*/0.8. This is because when the highest bid is the unique bid above the hard reserve *r* and below the soft reserve *r*/0.8, they win and pay their bid.

B. Dynamic Revenue Sharing v2

196. With DRSv2, AdX dynamically adjusted its take rate to sometimes be higher or lower than 20% (while maintaining the average take rate at 20% across auctions cleared by AdX) to win impressions that it would not have if the take rate was kept at 20% on a per-query basis.[292] In contrast, under DRSv1, AdX only adjusted take rates to be below 20%.[293] AdX decided to increase or decrease the take rate from 20% depending on a few factors such as the comparisons

[292] GOOG-NE-13207241 at -1. "AdX Dynamic Revshare v2: Launch Doc." [redacted]

[293] GOOG-TEX-00777528. Email thread, "Subject: Re: [Monetization-pm] Re: [drx-pm] LAUNCHED! AdX Dynamic Revenue Share (DRS)." ("We lower the revenue share per query as needed.")

between the first and the second highest bid and the publisher reserve, as well as debt balances for the publishers and advertisers.[294]

197. DRSv2 was launched in the second half of 2016.[295] Google announced DRSv2 when it was launched.[296] The publishers were allowed to opt out of DRSv2, however, if they did, Google turned off DRSv1 for these publishers as well.[297] Advertisers and ad buying tools could not opt out of DRSv2. [298]

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[295] GOOG-NE-04934281 at -81. July 30, 2018. "Dynamic Revenue Share." (describing release dates of "feature flags" (functionality), including that it launched into AdX UI in June 2016 and came into effect in August 2016.)

[296] GOOG-NE-06842715 at -20. May 10, 2016. "AdX Auction Optimizations." (describing that DRS would be announced in June 2016.)

[297] GOOG-NE-04934281 at -86. July 30, 2018. "Dynamic Revenue Share." ("You may choose to opt-out of revenue share based optimizations in the AdX UI. If you opt-out we will apply your contracted revenue share to every Open Auction query and you will not benefit from the increased revenue from this optimization.")

[298] GOOG-NE-04934281 at -85. July 30, 2018. "Dynamic Revenue Share." ("Q: Can buyers opt out? A: Revenue share based optimizations are controlled by sellers only.")

[299] For the sake of clarity, I abstract away from the ad buying tool fee, since it is immaterial to the analysis conducted here.

[redacted]

[redacted]

199. Under DRSv2, AdX is allowed to either increase the take rate on a per impression basis or decrease it, as it can be seen from its description. This is one important difference in comparison to DRSv1 where AdX was only allowed to decrease the take rate. AdX decides to increase or decrease the take rate based on the average take rate it applied to that publisher in the billing period. If it is close enough to the contractual requirement of 20%, then AdX sometimes decreases the take rate to win impressions that it would not have won otherwise. If it is much lower than 20% and the top bid is high enough compared to the floor, AdX increases the take rate to recoup the losses it incurred in auctions where it decreased its take rate.

200. To illustrate how DRSv2 works, imagine an impression arrives, and the publisher reports a price floor of $10. AdX solicits bids and receives top two bids of $20 and $15. In this case, with both a regular second-price auction and DRSv1, the AdX clearing price is $15, because $15 is greater than $10/0.8 = $12.5. AdX takes a 20% take-rate of $3 and passes on $12 to the publisher. Under DRSv2, if the winning advertiser has some debt from prior impressions, DRSv2 then

[redacted].

[redacted]

additionally charges the advertiser $0.5, clears $0.5 of its debt, keeps a 20% take rate on this and passes $0.4 to the publisher. In addition, if the publisher has some debt from prior impressions, DRSv2 then additionally withholds $1 from the publisher and clears $1 of its debt. So altogether, the advertiser pays $15.5 instead of $15, and the publisher is paid $11.4 instead of $12. This example is illustrated in Figure 35 below.

Figure 35: An impression is cleared by AdX with the DRSv2 dynamic adjustment of the take rate. Advertiser pays more, publisher earns less due to debt clearing

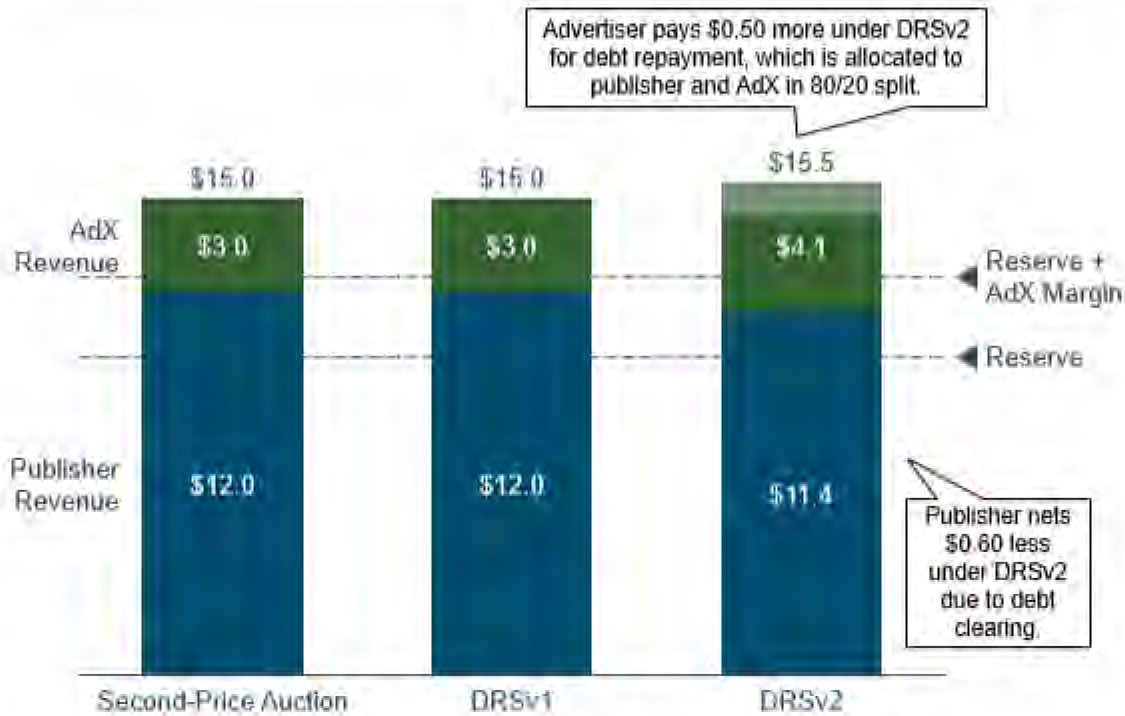

201. Imagine that another impression arrives, and the publisher again reports a price floor of $10. AdX solicits bids and this time receives top-two bids of $20 and $10. In this case, with both a regular second-price auction and DRSv1, the AdX clearing price is $12.5, because $20 is greater than $12.5 > $10. AdX takes a 20% take rate of $2.5 and passes on $10 to the publisher. Under DRSv2, if the winning advertiser has some debt from prior impressions, DRSv2 then additionally charges the advertiser $0.5, clear $0.5 of its debt, keeps a 20% take rate on this and passes $0.4 on to the publisher. In addition, if the publisher has some debt from prior impressions, DRSv2 then additionally withholds $0.2 from the publisher and clears $0.2 of its debt. So altogether, the advertiser pays $13 instead of $12.5, and the publisher is paid $10.2 instead of $10. This example is illustrated in Figure 36 below.

Figure 36: An impression is cleared by AdX with the DRSv2 dynamic adjustment of the take rate. Advertiser pays more, publisher earns more due to debt clearing

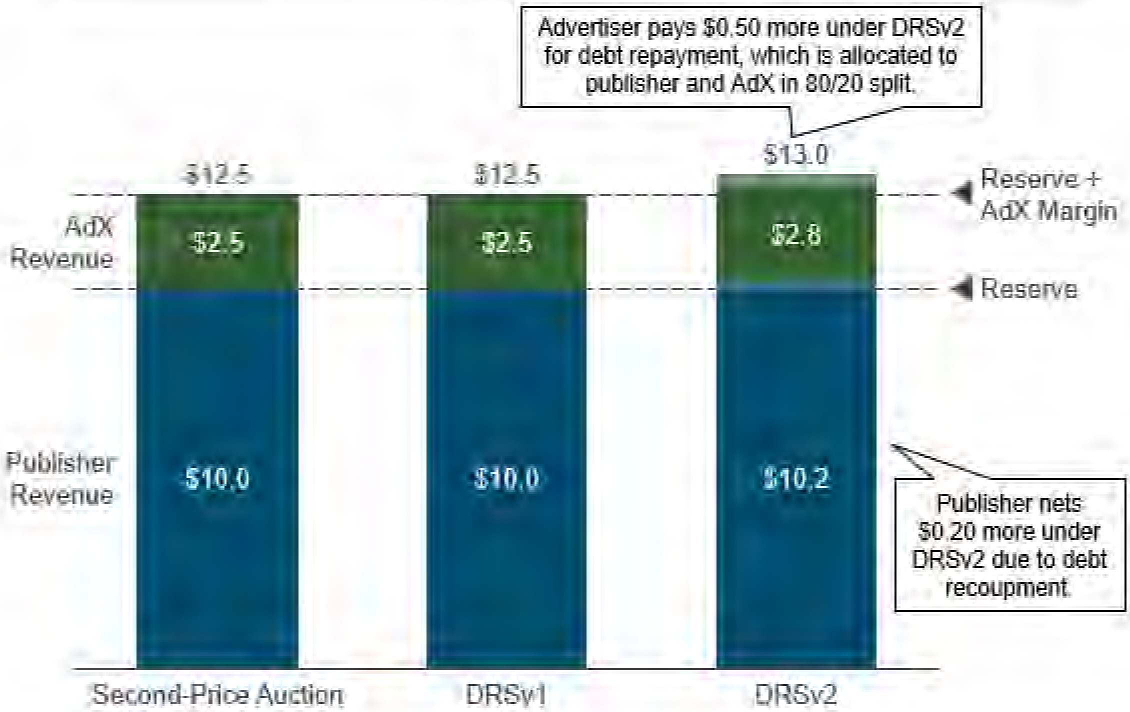

202. Imagine another impression arrives, and the publisher again reports a price floor of $10. AdX solicits bids and this time receives top two bids of $9 and $8. In this case, with both a regular second-price auction, DRSv1, or DRSv2, the impression is not transacted with AdX,[303] because even with a take-rate of 0% AdX still cannot clear the price floor. This example is illustrated in Figure 37 below.

[303] Note that the impression might be sold to another exchange or go unsold.

Figure 37: An impression that AdX could not clear even with the DRSv2 dynamic adjustment of the take rate

203. Imagine that a final new impression arrives, and the publisher again reports a price floor of $10. AdX solicits bids and this time receives top two bids of $11 and $10. In this case, with a regular second-price auction and a static take-rate of 20%, the impression would not be transacted with AdX.[304] Because that would require a top bid of at least $12.5 for a 20% take-rate to still result in a payment above the price floor of $10. With DRSv1 or DRSv2, provided that the average take-rate with this publisher during this billing period is sufficiently high,[305] DRSv1 or DRSv2 would instead clear the transaction, charge the winner $11, collect $1 and pass on $10 to the publisher for a fee of approximately 9.1%. DRSv2 further tracks that the advertiser paid $11 for an impression that 'costs' $12.5 and adds $1.5 to that advertiser's debt. DRSv2 also further tracks that the publisher was paid $10 on a total AdX revenue of $11, whereas 80% of $11 would have been $8.8, and therefore also adds $1.2 to the publisher's debt. If these four impressions constitute the same winning advertiser and publisher, then the advertiser has cleared $1 of debt in the first two impressions and accumulated $1.5 of debt in the final example, and so will later pay AdX to clear the remaining $0.5 of debt. The publisher has cleared $1.2 in the first two

[304] Note that the impression might be sold to another exchange or go unsold.
[305] If the average take-rate with this publisher is too high, then DRSv1 and DRSv2 both conclude as a regular second-price auction.

impressions and accumulated $1.2 of debt in the final example, and so is now balanced.[306] This example is illustrated in Figure 38 below.

Figure 38: An impression cleared by AdX with the DRSv2 dynamic adjustment of the take rate. Both publisher and advertiser accumulate debt

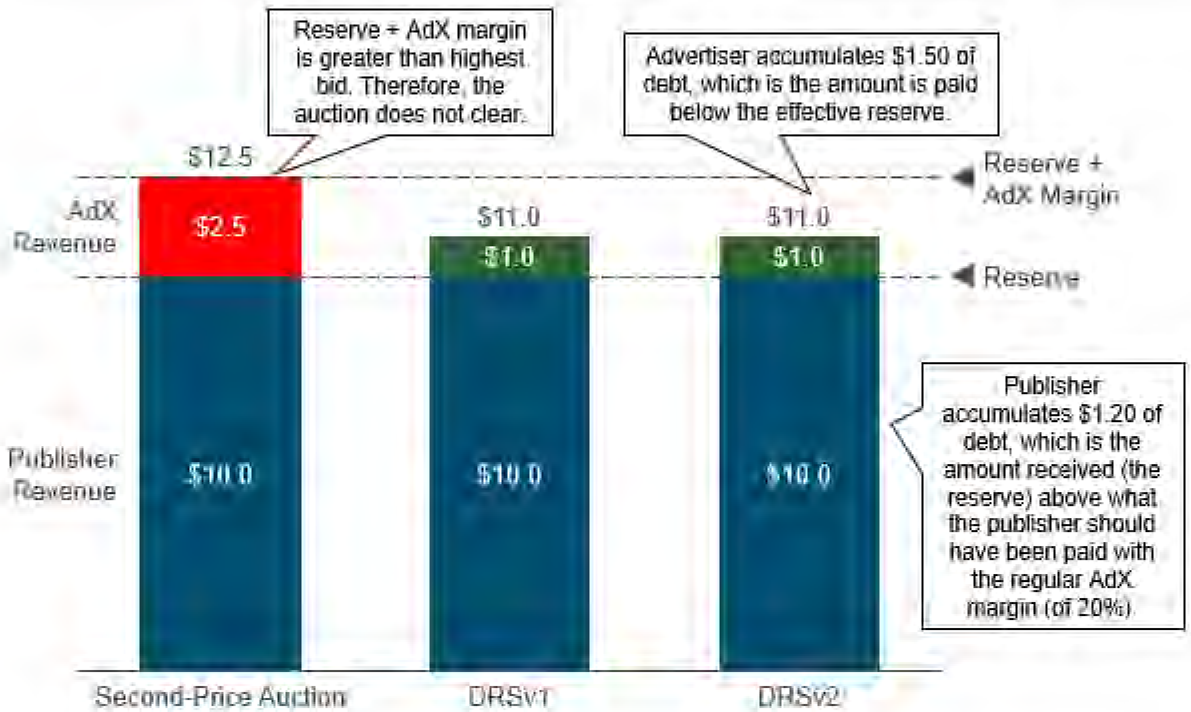

204. In these examples the advertiser wins the same impressions under DRSv1 and DRSv2 but pays more under DRSv2 (because DRSv2 tracks 'debt' incurred when DRSv1 activates and charges the advertiser to make up for this later).

a. As compared to no DRS, the advertiser wins an additional impression, but pays their bid for it under DRSv1 (and under DRSv2, the advertiser not only pays their bid for the impression but further accumulates debt they must pay later!).

[306] For this final impression, observe that the advertiser is paying their bid. In particular, if the advertiser had reported a bid of $12 instead, they would have been charged $12 (and AdX would have taken $2, for a fee of approximately 16.66%). If the advertiser had reported a bid of $10.1, they would have been charged $10.1 (and AdX would have taken $0.1, for a fee of approximately 1%). That is, while both DRSv1 and DRSv2 act like a truthful second-price auction for some bids, they are instead a first-price auction in the 'dynamic range' from $10 to $12.5, where AdX will charge the winner their bid.

Therefore, the advertiser achieves the same payoff under no DRS and DRSv1, and strictly lower utility under DRSv2 than either.[307]

b. As for the publisher, they sell the same impressions to AdX under DRSv1 and DRSv2 but receive $0.4 less payout under DRSv2 (because DRSv2 tracks 'debt' incurred when DRSv1 activates and lowers the take rate and pays less to the publisher later to make up for this). However, different circumstances could have resulted in DRSv2 paying out more than DRSv1 (for example, if the advertiser cleared all their debt in the first two impressions, the publisher would have received the same total payout under DRSv1 and DRSv2). In comparison to no DRS, the publisher receives an additional $10 from AdX across all impressions under DRSv1, but an unknown amount less from the counterfactual of selling the fourth impression elsewhere. If the publisher would have received more than $10 for this impression elsewhere, then no DRS yields greater revenues. If they would have received less than $10, then DRSv1 yields greater revenues. The comparison of DRSv2 to no DRS in this example is nearly identical, except that the publisher loses $0.4 on the first two impressions (although again, this aspect could have gone differently with different numbers). Finally, observe that because DRSv2 "clears publisher debt", DRSv2 can "do more DRSv1" (that is, over an extended period of time, there would also be impressions where DRSv1 acts like a clean second-price auction with take rate of 20% and fails to transact the impression through AdX, while DRSv2 instead clears it).

205. Since under DRSv2 the advertisers sometimes pay their bid, DRSv2 corresponds to a dirty second-price auction under some conditions and as a result, DRSv2 is not a truthful auction format. Specifically, whenever DRSv2 charges the winning advertiser their bid and pays the publisher their reserve, it is a dirty second-price auction with a hard reserve of r and a soft reserve of $r/0.8$. This is because, in the naming scheme of the dirty second-price auction definition I

[307] Note that the only change between no DRS and DRSv1 is that the advertiser now wins impressions and pays their value. For a sophisticated advertiser who has truly accounted for their budget constraints and knows a value so that they are truly indifferent between paying this value to win the impression or losing, they are truly no better or worse off. For a less sophisticated advertiser who has not fully incorporated budget constraints into their value, they may prefer not to win an impression and pay their value, because this hurts their 'return on investment (ROI)' – getting nothing might be preferable to investing $10 into advertising for a return of exactly $10.

provide above, when the highest bid b_1 is the unique bid above the hard floor *r* and below the soft floor *r*/0.8, they win and pay their bid.[308]

206. DRSv2 goes a step further than being a dirty second-price auction, due to its debt mechanism. In particular, when an advertiser submits a winning bid in the 'dynamic region' (*i.e.,* a bid between the publisher-set reserve of *r* and the effective reserve of *r*/0.8), the advertiser not only pays their bid now resulting in a payoff of 0, but further accumulates debt that must be paid later. That is, *assuming that all debt is eventually cleared,* the auction that advertisers participate in under DRSv2 is as described below, which I'll call a *debt-aware second-price auction*.

a. AdX learns its reserve *r* from the ad server and solicits bids from the ad buying tools. Let b_1 denote the highest received bid and b_2 denote the second highest.[309]

b. If the second highest bid is above the effective reserve (*i.e.*, $b_2 \geq r/0.8$), then the clearing price is b_2.

c. If the effective reserve is between the highest and second highest bids (*i.e.*, $b_1 \geq r/0.8 > b_2$), then the clearing price is *r*/0.8.

d. If the highest bid is high enough to clear the reserve, but not the effective reserve (*i.e.*, $r/0.8 > b_1 \geq r$), one of the following happens:

   i. AdX fails to win the impression. This happens during periods when AdX's average take-rate during the current billing period with this publisher is much lower than 20%.

   ii. The clearing price is *r*/0.8, the effective reserve, which strictly exceeds the winning bidder's bid.[310]

e. If the highest bid is below the reserve (*i.e.*, $b_1 < r$), no bid is returned and AdX does not win the impression.

---

[308] This is the only conclusion that does not hold for all three payment variants articulated in GOOG-NE-13207241. All of these variants are still not truthful, but do not verbatim fit the 'dirty second-price auction' framework. GOOG-NE-13207241. "AdX Dynamic Revshare v2: Launch Doc." (under the section "Other Pricing Rules")
[309] For the sake of clarity, I abstract away from the ad buying tool fee, since it is immaterial to the analysis conducted here.
[310] This follows because, assuming all debt clears, the advertiser accumulates debt equal to *r*/0.8 minus their payment, which totals *r*/0.8 of debt plus payment.

207. I discuss the implications of participating in a debt-aware second-price auction later in this section, and observe now just that step d.ii. above indeed charges the winning bidder a price that exceeds their bid.[311]

C. Truthful Dynamic Revenue Sharing

208. With the third and last iteration tDRS, AdX dynamically adjusted its take rate to sometimes be higher or lower than 20% to win impressions that it would not have if the take rate was kept at 20% on a per-query basis. Under tDRS, Google determined the dynamic take rate it is going to charge before 'peeking' at the bids.[312] In contrast, both DRSv1 and DRSv2 adjusted the take rate after AdX observed the submitted bids.[313] Under tDRS, the take rate calculation is done based on the past AdX data. Internal Google documents states that the prediction model that determines the take rate "predicts for a given query whether a specific buyer would bid above the pre-revshare reserve price."[314]

209. tDRS was fully launched in the second half of 2018.[315] When AdX migrated to a first-price auction format in 2019, the DRS program was shut off.[316]

[redacted]

- [redacted]
- [redacted]

[311] And recall again that this follows because we are assuming all debt clears, and so therefore can count payments made to clear debt the moment that debt is accumulated.
[312] GOOG-AT-MDL-019244499. "Truthful DRS Auction Walkthrough." ("For each buyer, its reserve price revshare factor will be determined based prediction result before the request is being passed down to RTBs or CAT2 mixer (for Adwords and DBM).")
[313] GOOG-NE-13226622 at -2. "Truthful DRS Design Doc." ("One known issue with the current DRS is that it makes the auction untruthful as we determine the AdX revshare after seeing buyers' bids and use winner's bid to price itself (first-pricing) when the bid is within the dynamic region.")
[314] GOOG-NE-13214748 at -8. "Modeling Design Doc for Truthful DRS."
[315] GOOG-TEX-00858434. January 29, 2020. "Dynamic Revenue Share." ("Update (July 30, 2018): We launched a new DRS model (tDRS).")
[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

211. Observe that from the advertiser perspective, tDRS is always a regular second-price auction with reserve r*. As a result, tDRS is truthful, which is different from DRSv1 and DRSv2. This indeed was Google's motivation behind creating tDRS, stating that "the current DRS … makes the auction untruthful as we determine the AdX revshare after seeing the buyers' bids and use the winner's bid to price itself (first-pricing) when the bid is in the dynamic region. This could incentivize buyers to bid strategically instead of truthfully."[325]

212. To illustrate how tDRS works, imagine an impression arrives, and the publisher reports a price floor of $10. AdX first decides on an alternate reserve of $9 and a beat-the-floor revenue share of 80%, without seeing any bids. From both the advertisers' and publisher's perspective this executes as a clean second-price auction with a reserve of $12.5 and a take-rate of 20%.

213. Imagine another impression arrives, and the publisher again reports a price floor of $10. AdX again decides on an alternate reserve of $9, but this time a beat-the-floor revenue share of 100%, without seeing any bids. From the advertisers' perspective, this is a second-price auction with reserve $11.25.[326] From the publisher's perspective, their payout is complicated. The auction will clear whenever b_1 exceeds $11.25. In this case, if b_2 happens to exceed $12.5, then the publisher will receive 80% of b_2, a clean 20% take-rate. If instead b_2 falls below $12.5, then the publisher will receive $10, their price floor, and necessarily have a take-rate of < 20% (because both b_2 and the reserve are less than $12.5, so the winning bidder pays less than $12.5). If for

[redacted]

[325] GOOG-NE-13226622. "Truthful DRS Design Doc."
[326] This follows as $9/0.8 = $11.25, which exceeds $10.

example b_2 is $12, then the publisher receives a revenue share of $10/$12 for roughly 83.3% > 80%. An 80% revenue share would have given $9.60, so a debt of $0.4 is accumulated.[327] If b_2 is less than $11.25, then the publisher receives a revenue share of $10/$11.25 for roughly 88.9% > 80%. An 80% revenue share would have given $9, so a debt of $1 is accumulated.[328]

214. Imagine a final impression arrives, and the publisher again sets a price floor of $10. This time, AdX decides on an alternate reserve of $15 (because AdX believes the impression can fetch bids much higher than $10), and a beat-the-floor revenue share of 100%.[329] From the advertisers' perspective, this is a second-price auction with reserve $18.75. From the publisher's perspective, their payout is complicated. The auction will clear whenever b_1 exceeds $18.75. In this case, if b_2 happens to exceed $18.75, then the publisher gets 80% of b_2 for a clean 20% take-rate. If b_2 instead falls below $18.75, then the 80% payout to the publisher would be $15, but AdX uses this opportunity to clear some of the publisher's 'debt'. Instead of paying the publisher $15, AdX pays out $15 less $y and clears $y 'debt'. $y will be at most $5 (to guarantee that the publisher sees a payment exceeding their price floor of $10), and also at most $15-0.8*$b_2$ (to guarantee that the publisher sees a payment exceeding 80% of b_2).

215. Assuming that all debt is cleared, the publisher receives an 80% revenue share across all auctions AdX clears, and AdX receives payment exactly according to a second-price auction with reserve $r^* = \max\{R/0.8, r/x\}$. In particular, if we count debt when it is accumulated rather than when it is ultimately paid (call this the publisher's debt-aware payout), then the publisher receives 80% revenue share of a second-price auction with reserve $r^* = \max\{R/0.8, r/x\}$ on a per-auction basis. In particular, observe that if $r > R$ and $x > 0.8$, *the publisher's debt-aware payout can be less than their price floor*. In my previous example, when the publisher reports a price floor of $10, and AdX runs a second-price auction with reserve r^* = $11.25, and b_2 < $11.25 < b_1, the impression clears at a price of $11.25, and therefore the publisher's debt-aware payout is 80% of $11.25 = $9 < $10.[330]

[redacted]

[329] When the alternate reserve exceeds the publisher-set reserve, the beat-the-floor revenue share is immaterial.

[330] If tDRS calculated debt according to the formula $r^*(x-0.8)$ instead of $r-\max\{0.8^*r/x, 0.8^*b_2, R\}$, and only used values of x = 1.0 or x = 0.8, then tDRS overcalculates debt, and would therefore eventually pay out publishers even less than I compute above. That is, AdX would collect revenue according to a second-price auction with reserve $r^* = \max\{R/0.8, r/x\}$, while giving the publisher at most an 80% debt-aware revenue share, and sometimes strictly less.

D. Impact of DRS Programs on Publishers

216. The impact of DRSv1 on the publishers, as compared to no DRS, has an indeterminate impact on publisher revenue. In a situation where DRSv1 kicks in and helps an impression clear when it would not have otherwise, the publisher is always paid their price floor. This price floor could be greater than, equal to, or less than the revenue the publisher would have gotten had AdX not transacted the impression. For example:

a. The price floor might be equal to the highest bid received via header bidding. In this case, the publisher would receive exactly the price floor in case AdX does not transact this impression (because the price floor is a bid the publisher has in hand). Therefore, in this case, the publisher is neutral towards DRSv1.

b. The price floor might be a reserve set on exchanges by the publisher, and it could be that no exchanges cleared this reserve in header bidding. In this case, the outside opportunity if AdX does not transact this impression is zero, because it would otherwise go unsold. In this case, the publisher sees increased revenue from DRSv1.

c. The price floor might be set via (Enhanced) Dynamic Allocation with other exchanges participating via the waterfall. In this case, the publisher does not know exactly what opportunity would have arisen had AdX not transacted the impression and it instead entered the waterfall. It could be that the publisher would have gotten lucky and seen revenue from the waterfall that exceeds this price floor. In this case, the publisher sees decreased revenue from DRSv1. It could be that the publisher would have gotten unlucky, the impression would not have cleared in the waterfall. In this case, the publisher sees increased revenue from DRSv1. On average, if the publisher sets a reserve strictly above their expected opportunity cost, they would expect to see increased revenue from DRSv1 (although this is not guaranteed to occur impression-by-impression).

217. DRSv2 also has an indeterminate impact on publisher revenue, as compared either to no DRS or DRSv1. DRSv2 has two channels of additional effects, both of which have indeterminate impact: DRSv2 could clear more or fewer transactions in the dynamic region as compared to DRSv1. If advertisers bid truthfully in DRSv2 and DRSv1, then DRSv2 would clear more transactions in the dynamic region, because DRSv2 increases its take-rate on some transactions

so that it can lower its take rate in the dynamic region more frequently. If some advertisers were to realize they are participating in a debt-aware second-price auction and stopped bidding in the dynamic region under DRSv2, then fewer transactions would be cleared in the dynamic region.[331] Therefore, if a publisher saw increased revenue from DRSv1 to no DRS, they would also see increased revenue from this aspect of DRSv2 to no DRS. If a publisher saw decreased revenue from DRSv1 to no DRS, they would also see decreased revenue from this aspect of DRSv2 to no DRS. Whether a publisher saw increased revenue from this aspect of DRSv2 to DRSv1 could go either way, depending both on whether they saw increased or decreased revenue from DRSv1 to no DRS, and whether more or fewer impressions transact in the dynamic region.

218. DRSv2 maintains a single debt tracker for each advertiser and each publisher, rather than each (advertiser, publisher) pair.[332] This means that an advertiser can incur debt while purchasing an impression from one publisher (and the publisher incurs 80% of that debt) and clear it while purchasing an impression from another publisher (and this publisher then earns 80% of the cleared debt). This aspect is zero-sum across publishers but causes transfers from some publishers to others. A publisher could see increased or decreased revenues due to this aspect, in comparison to both DRSv1 and no DRS, depending on whether they tend to participate more in debt-building (decreases revenue) or debt-clearing (increases revenue) transactions.

1) Publishers would have set different reserve prices to maximize their revenues had Google revealed DRSv1

219. Google concealed material information from publishers by not disclosing the implementation of DRSv1. Since publishers believed that AdX runs a regular second-price with their given reserve price and a static take rate of 20%, a strategic publisher would set a price floor that maximized their revenue under these circumstances. Had they known AdX dynamically adjusted its take rate, publishers would set different price floors.[333] In general, auction formats are known to be vital for optimizing revenues, as are reserve prices. As a publisher, it is therefore

[331] I discuss advertiser incentives under DRSv2 in Section VII.B. In particular, bid-shading in the dynamic region results in the same outcomes as truthful bidding. Outcomes change only if advertisers forego the dynamic region entirely.
[332] *See* GOOG-NE-13207241 at -45. "AdX Dynamic Revshare v2: Launch Doc." [redacted]
The example at -46 demonstrates advertiser b_1 incurring debt with publisher p_1 and clearing it with publisher p_2.
[333] For example, with a static take rate of 20%, a publisher might naturally set the price floor r such that r/0.8 is the revenue-optimal reserve, taking into account the opportunity cost of selling the impression elsewhere. If the publisher knew that the take-rate could go down [redacted] on average, a publisher might naturally set the price floor r such that [redacted] revenue-optimal reserve (again, taking into account the opportunity cost of selling the impression elsewhere). That is, one natural publisher response to DRSv1 is to increase price floors on AdX.

material to understand what auction AdX is running and how the reserve set on AdX influences the advertiser bids.[334]

E. Impact of DRS Programs on Exchanges

220. DRSv1 would lead to an increase in AdX's win rate and an increase in AdX's revenue as compared to no DRS. Since DRSv1 was not disclosed to the advertisers, they would still bid their true values for the impression. The only change that happens with DRSv1 is that sometimes AdX successfully clears an impression when it would have not without DRSv1. Therefore, AdX would win at least every impression that it wins without DRSv1 while paying the same price and might win additional impressions. This leads to an increase its win rate and revenue. This holds whether other exchanges participate via waterfalling or header bidding.[335]

221. Internal Google documents show that AdX's win rate and revenue indeed increased as a result of DRSv1. An internal email provides the immediate impact of DRSv1 a week after the launch, stating that it brought an additional (annualized) [redacted] in revenue to AdX.[336] Furthermore, the email states "overall match rate for AdX publishers increases by [redacted] when selling to AdX buyers,"[337] demonstrating the increase in the number of transactions cleared by AdX. An internal Google presentation stated that DRSv1 led to [redacted] increase in annual recurring revenue.[338]

222. DRSv2 would lead to an increase in AdX's win rate and revenue as compared to no DRS. The impact of DRSv2 on AdX win rate and revenue is indeterminate as compared to DRSv1. Because DRSv2, via its debt mechanism, enables AdX to dynamically adjust its take rate more often compared to DRSv1, this aspect would lead to an increase in win rate for AdX compared to DRSv1. Since DRSv1 leads to an increase in AdX win rate compared to no DRS, this means that DRSv2 would lead to a further increase in win rate over no DRS as well. However, since DRSv2

---

[334] For example, because DRSv1 is material to advertisers' bid decisions, if DRSv1 were to be abruptly disclosed, it could cause advertisers to abruptly start bid shading in AdX, which could abruptly negatively impact a publisher's revenue. As another example, perhaps the publisher is sophisticated and deciding how much to adjust Value CPMs to mitigate AdX's Last Look advantage, which was exacerbated by DRSv1. In such a situation, the publisher might underestimate the impact of AdX's Last Look advantage and mistakenly decide against adjusting Value CPMs further.
[335] These additional impressions might have otherwise been cleared by a non-Google exchange, and therefore that exchange's win rate and revenue would decrease. It is also possible that some of these additional impressions would have otherwise gone unsold, which would not directly impact other exchanges.
[336] GOOG-TEX-00777528. Email thread, "Subject: Re: [Monetization-pm] Re: [drx-pm] LAUNCHED! AdX Dynamic Revenue Share (DRS)."
[337] GOOG-TEX-00777528. Email thread, "Subject: Re: [Monetization-pm] Re: [drx-pm] LAUNCHED! AdX Dynamic Revenue Share (DRS)."
[338] GOOG-NE-06842715 at -18. May 10, 2016. "AdX Auction Optimizations."

is not truthful and aspects of it were disclosed, advertisers may have shaded their bids. I have previously noted that bid-shading itself within the dynamic region itself does not change outcomes under DRSv2, but if some advertisers choose to skip the dynamic region entirely, this could cause fewer impressions to transact in the dynamic region compared to DRSv1. Still, if any advertisers transact in the dynamic region, DRSv2 clears additional impressions beyond no DRS, and therefore DRSv2 leads to increased win rate and revenue for AdX in comparison to no DRS.[339] When compared to DRSv1, AdX's win rate and revenue could increase or decrease, depending on the magnitude of the two effects noted above which push in opposite directions.

223. tDRS would lead to an increase in AdX's revenue and win rate as compared to no DRS. tDRS and no DRS are both truthful auctions, so advertisers would submit the same bids to both. The only distinction between tDRS and DRS is that tDRS gives AdX more flexibility over the effective reserve for its auction (under no DRS, it must be at least r/0.8, where r is the publisher's price floor,[340] with tDRS, AdX can now set the effective reserve as low as r). AdX would therefore use this flexibility to better optimize its revenue. Moreover, any optimization would come by lowering the effective reserve, which would increase AdX's win rate. Therefore, tDRS would lead to both an increase in AdX's revenue and win rate.[341]

224. DRSv1, DRSv2, and tDRS all naturally exacerbate all conclusions regarding non-AdX exchanges under Dynamic Allocation and Enhanced Dynamic Allocation. In Section IV, I established that exchanges without AdX's advantage under Dynamic Allocation and Enhanced Dynamic Allocation clear fewer impressions and earn less revenue compared to AdX. The gap in win rate and revenue between exchanges with and without AdX's Last Look advantage would be larger when the Last Look advantaged exchange uses some form of DRS in comparison to when all exchanges use fixed take rates. The "Last Look" advantage bestowed upon AdX via Dynamic Allocation and Enhanced Dynamic Allocation gives AdX a reserve price such that if AdX can pay more than this reserve, AdX is guaranteed to win. When other exchanges participate in the waterfall, Dynamic Allocation prevents other exchanges from having an opportunity to return live bids. For Dynamic Allocation with other exchanges in header bidding, this allows AdX to use the

[339] Again, these additional impressions might have otherwise been cleared by a non-Google exchange, and therefore that exchange's win rate and revenue would decrease. It is also possible that some of these additional impressions would have otherwise gone unsold, which would not directly impact other exchanges.
[340] Note that a sophisticated publisher could respond to tDRS by increasing their price floors, which would limit AdX's flexibility.
[341] Again, these additional impressions might have otherwise been cleared by a non-Google exchange, and therefore that exchange's win rate and revenue would decrease. It is also possible that some of these additional impressions would have otherwise gone unsold, which would not directly impact other exchanges.

maximum live bid to inform its reserve. For Enhanced Dynamic Allocation, this allows AdX to use direct deal CPMs as its reserve price. DRSv1, DRSv2 and tDRS give AdX more flexibility to clear this reserve price determined by Dynamic Allocation or Enhanced Dynamic Allocation more often, which further increases Google's win rate and revenue and further decreases the win rate and revenue of other exchanges.[342]

F. Impact of DRS Programs on Advertisers

225. DRSv1 was likely neutral to AdX advertisers' payoffs, although may have had some negative impact. First, DRSv1 was concealed, and therefore AdX advertisers would bid their true values into what they believed was a second-price auction. Because advertisers pay their bids when winning an impression in the dynamic region, their payoff is 0 when this occurs (the same as if they did not win). Still, there are two mechanisms by which AdX advertisers may have been negatively impacted. First, if publishers were to regularly re-optimize reserves, their reserves might have naturally increased under DRSv1 (even without fully understanding AdX's auction). If publishers' reserves increase, AdX advertisers pay more for impressions, and therefore suffer negative impact. Second, some advertisers might care about ROI in addition to payoff. Winning an impression that gives payoff 0 harms ROI, and therefore such advertisers would also suffer a negative impact.

1) DRSv2 led to a decrease in AdX advertisers' payoffs

226. DRSv2 was quite negative towards AdX advertisers' payoffs.[343] I previously noted that from AdX advertisers' perspective, DRSv2 is a debt-aware second-price auction. In particular, a debt-aware second-price auction is exactly a second-price auction with reserve r/0.8 *except if the winning bid is between r and r/0.8 it is treated as r/0.8 instead*. This means three things for how an advertiser with value v should optimally respond: (a) if v is outside the dynamic region (*i.e.,* v > r/0.8), the advertiser should optimally bid v, (b) if v is inside the dynamic region, shading v to some other value inside the dynamic region makes no difference,[344] (c) if v is inside the dynamic region, *the advertiser is best off omitting a bid entirely*. Let me now draw a few conclusions.

[342] Note that I present this conclusion in Section IV on Dynamic Allocation and Enhanced Dynamic Allocation as well.
[343] At least, those who used third-party ad buying tools, as DV360 and GDN were exempted from DRSv2. GOOG-TEX-00831090 at -1. April 17, 2017. "DRX 2.0 Quality." ("Sellside dynamic revenue share (DRS) is applied to AdX RTB but not Adwords or DBM.")
[344] Conditional on this advertiser still winning, bid shading within the dynamic region benefits advertisers only because it makes the advertiser less likely to win in the first place but does not change the outcome if the advertiser still wins. Of course, depending on the extent to which advertisers fully understood the precise mechanics of DRSv2, they may have shaded their bids instead of skipping the dynamic region (depending on how well advertisers fully understood debt repayment, they may have even shaded their bids outside of the dynamic region since bid shading

a. If all advertisers responded optimally to DRSv2, *no advertiser would bid in the dynamic region, and therefore DRSv2 would be equivalent to no DRS*. That is, exactly the same advertisers would win exactly the same impressions and pay exactly the same amount and the entire DRSv2 program would be obviated. Internal Google documents show that [redacted][345] This increase in AdX revenue is only possible if transactions are cleared in the dynamic region, and every transaction cleared in the dynamic region necessarily involves an advertiser paying more than their value for an impression, and therefore suffering decreased payoff in comparison to no DRS. I again wish to emphasize that any increase in AdX revenue *must* come at the cost of some advertisers having lower payoff under DRSv2 compared to no DRS, *and* those advertisers behaving sub-optimally under DRSv2.

b. No matter how an advertiser responds to DRSv2, they cannot possibly be better off than no DRS. In particular, the best they can do is to bid truthfully outside the dynamic region and avoid the dynamic region entirely. Doing so will give them the exact same outcomes as under no DRS, and any other behavior causes a decrease in their payoff as compared to no DRS.

c. Bid-shading in the dynamic region does not improve an advertiser's outcomes over truthful bidding,[346] they must skip the dynamic region entirely in order to better respond to DRSv2 over truthful bidding. This fact only follows from the precise mechanism that DRSv2 uses to track and clear debt and cannot be deduced merely from the fact that DRSv2 dynamically increases and decreases AdX's take rate.

outside the dynamic region does cause advertisers to pay less when they win, but this is only because they are clearing less debt). Moreover, if advertisers did not know for which auctions DRSv2 was active, advertisers who considered bid shading may have further shaded their bids even on AdX auctions that concluded as truthful second-price auctions without DRSv2.

[345] GOOG-NE-13234466 at -67. "Overall Pub Yield With DRS(v2)."

[346] Again, bid shading could only help by reducing the risk that the advertiser wins, or causing the advertiser to transact in the dynamic region less often. Conditioned on winning the same set of impressions, bid shading has no impact.

2) Advertisers would have submitted different bids to maximize their payoffs had Google revealed DRSv1

227. Google misled advertisers by not revealing DRSv1, and hence led them to believe the AdX auction was a regular second-price auction, which would cause them to engage in suboptimal behavior. When advertisers believe they are participating in a regular second-price auction, they would bid their true value for the impression, because it is a truthful auction. However, DRSv1 is not truthful, as established before. Therefore, concealing DRSv1 caused advertisers to bid their true value in a non-truthful auction, whereas advertisers would get higher a higher gain by shading their bids.

228. By not revealing DRSv1 to the advertisers, Google made material gains. This is because if advertisers were to shade their bids, which is the natural bidding behavior in a non-truthful auction like DRSv1, this would lead to less revenue for both AdX and publishers.[347] However, advertisers likely did not shade their bids, since Google never publicly revealed DRSv1.

G. Some aspects of DRS are exceptionally misleading

229. To conclude the section, I want to briefly note a few aspects of DRSv2 that I find exceptionally misleading to advertisers, and an aspect of DRSv2 and tDRS that I find misleading to publishers. Much of my analysis below concerns the concept of 'debt' to mislead both advertisers and publishers regarding how much they are paying or paid out.

230. First, I want to repeat that my previous analysis establishes that when advertisers behave optimally in DRSv2, no transactions should ever occur in the dynamic region. Instead, Google claims that enough transactions occurred in the dynamic region to account for [redacted].[348] This increase necessarily comes at the expense of advertisers ultimately *paying more than their value for an impression*.

231. Next, I want to highlight aspects of Google's description[349] regarding DRS that I find misleading, due to the concept of debt.

---

[347] But this does not mean DRSv1 as a whole, after accounting for bid shading, would necessarily lose revenue. It merely means that when comparing "DRSv1 when advertisers shade their bids" yields higher payoff for advertisers and less revenue for AdX and the publishers than "DRSv1 when advertisers bid their true values," which motivates concealing DRSv1.
[348] GOOG-NE-13234466 at -67. "Overall Pub Yield With DRS(v2)."
[349] GOOG-NE-04934281 at -84. July 30, 2018. "Dynamic Revenue Share."

a. For DRSv2, Google states: "Buyers are never charged more than their bid," I find this claim exceptionally misleading to advertisers. If by 'charged', one means 'charged on this impression as immediate payment, ignoring any debt that will be paid later', then the sentence is technically true. But if by 'charged', one means 'charged on this impression either as immediate payment or as debt to be collected later', then the sentence is false. Indeed, any winner in the dynamic region is ultimately charged more than their bid after accounting for both immediate payment and debt to be collected later.

b. For DRSv2, Google states: "sellers are always paid at least their reserve." I find this claim misleading to publishers. If by 'paid', one means 'paid on this impression as immediate payment, ignoring any debt generated that must be paid back later', then the sentence is technically true. But if by 'paid', one means 'paid on this impression as immediate payment, after subtracting any debt assigned that will be collected later', then the sentence is sometimes false. On impressions transacted in the dynamic region, publishers are indeed paid their reserve as immediate payment, but are also assigned non-negative debt. This debt may wind up being cancelled (if the winning advertiser clears their own debt while later transacting with this publisher), owed (if the winning advertiser clears their own debt elsewhere), or yielding extra payout (if some other advertiser later clears debt incurred elsewhere with this publisher). Even after accounting for debt, it is plausible to say that "sellers as a whole are paid at least the reserve set on any cleared impression." But it is inaccurate to say "sellers are *always* paid at least their reserve," because some sellers are not ultimately paid their reserve.[350]

c. If the concept of debt was not clearly disclosed, the general description of DRS as per-query revenue share optimization is insufficient for advertisers to draw conclusions at the level I have drawn in my report. Moreover, even for advertisers who are already optimizing bids at a per-impression level, the concept of debt significantly obscures feedback. Indeed, a typical optimizer might ask questions of the form "if I change my bid on this impression, what change does that cause in my payoff from this impression?" The concept of debt now means that changing a

[350] Sellers would reasonably care about whether they are paid their reserve or not, as this reserve constitutes the minimum amount they've decided to accept in order to forego the opportunity cost of selling the impression elsewhere.

bidder's bid on an impression *might cause them to pay more on another impression* (when that debt is cleared). In order for advertisers to have sufficient information regarding DRSv2 in order to avoid paying more than their value for an impression, Google would have needed to disclose a somewhat precise description of the debt concept. I do not know whether Google indeed made such a disclosure, nor how it was made, but in my opinion such information is vital to advertisers, even if they were already aware in a general sense that DRSv2 optimizes revenue shares on a per-query basis.

d. In a Google communications document,[351] that seems to have been active after tDRS launched,[352] Google states "Before or after the July change, we still do NOT pay publishers below publisher's floor…" I find this claim quite misleading to publishers concerning tDRS. If by 'pay', one means 'pay as immediate payment, ignoring any debt generated that must be paid back later', then the sentence is technically true. But if by 'pay', one means 'pay on this impression as immediate payment, after subtracting any debt assigned that will be collected later', then the sentence is false. Indeed, after accounting for debt assigned, I previously noted that the debt-aware payment to the publisher is sometimes less than their price floor. Even after accounting for debt, it is fair to say that "any publisher can guarantee a debt-aware payment at least as high as a desired price floor r by setting AdX's price floor to 1.25*r." But I find it misleading to blanketly assert that the publisher is not paid below their floor without further nuance.[353]

VIII. CONDUCT ANALYSIS: PROJECT BERNANKE

232. In this section, I analyze the conduct referred to as Project Bernanke, which was later expanded to Project Global Bernanke, and Project First Price Bernanke (also called "The Alchemist"). I also draw conclusions regarding their effects. I show that that Projects Bernanke and Global Bernanke did not affect GDN advertisers' payoffs and could increase some publishers' revenues while decreasing others, however they led to a lower win rate for non-GDN ad buying tools and advertisers that use those ad buying tools. Furthermore, Project Bernanke and its variants led to an increased win rate for GDN buyers (and in the case of Projects Bernanke and

[351] GOOG-TEX-00858434 at -40. January 29, 2020. "Dynamic Revenue Share."
[352] The document seems to be updated through at least December 2019, which is after the tDRS launch.
[353] Sellers would reasonably care about whether they are paid their floor or not, as this floor constitutes the minimum amount they have decided to accept in order to forego the opportunity cost of selling the impression elsewhere.

Global Bernanke, without improving GDN buyers' utilities), which leads to an increased win rate and revenue for GDN (again in the case of Projects Bernanke and Global Bernanke, without assisting GDN advertisers at all).

233. Project Bernanke and all its variants can be understood as simultaneously facilitating the effects of collusion among GDN advertisers, without their knowledge, and overbidding in auctions. I explain this view in detail in this section. The starting point for Project Bernanke is Google's observation that [redacted].[354] Because GDN is 'second-pricing itself', GDN would benefit by lowering the second-highest bid it sends in order to lower the payment GDN must make to win the impression.[355] In isolation, this would be a pure transfer of funds from exchange/publisher to GDN, but would technically result in a high take rate by GDN towards its advertisers, compared to what is contracted. This aspect of Project Bernanke is akin to facilitating collusion among GDN advertisers (and in this case, without their knowledge).[356] The second half of Project Bernanke uses the savings from the first half and spends it to subsidize overbidding.[357] That is, the second half of Project Bernanke boosts the bids of its advertisers before sending them to AdX, but uses the funds from the first half to cover any payment made above the advertiser's true bid. In all Project Bernanke variants, the two halves balance out to (a) generate increased revenue and increased win rate for GDN, (b) balance GDN's take rate at the intended 14%,[358] (c) have an indeterminate effect on publisher revenue (the first half decreases publisher revenue while the second half increases it), (d) in the case of Project Bernanke and Project Global Bernanke, not improve GDN advertisers' payoffs at all. Below I describe Project Bernanke in greater detail, and Appendix H contains additional discussion on overbidding and collusion in first- and second-price auctions.

[redacted]

[355] GOOG-AT-MDL-001412616 at -19. [redacted]

[356] By using the word 'collusion', I do not mean to imply that it is 'wrong' from a pure auction theory perspective for a group of bidders to get together and jointly strategize on how to collectively bid, nor for an ad buying tool to facilitate this. My understanding is that other ad buying tools may have dropped their second-highest bid entirely. GOOG-NE-13200831 at -1. "The case for encouraging buyers to declare two bids." ("Currently, the only [AdX] buyer who is employing this strategy [sending two bids] is GDN.")
Still, GDN is indeed facilitating collusion by implementing a joint strategy for its advertiser pool together, rather than processing each advertiser's bid in isolation.

[357] GOOG-AT-MDL-001412616 at -20. "Project Bernanke and margins story." ("What if overbid? We could bid too much. But we have to subsidize it. One is good for us [GDN] and bad for publishers. Other is bad for us [GDN] and good for publishers.")

[358] Some documents I have reviewed states that the GDN take rate is 14% and others state that it was 15%. I use 14% throughout the text, except in the cases where I cite a specific document that states 15%. This difference in the take rate does not have any effect on my conclusions throughout this section.

A. Project Bernanke

234. Prior to any modifications, GDN ran an internal auction (called the "CAT2" auction)[359] with only GDN advertisers and submitted the top two bids from that auction to AdX.[360, 361] Under **Project Bernanke**, between 2013 and 2015,[362] GDN manipulated advertisers' bids before sending them to AdX in the following manner:[363]

a. [redacted]

[redacted]

[redacted]

[redacted]

[359] Also called the "CAT2 auction." GOOG-NE-11753797 at -37. February 11, 2019. "DVAA Quality, Formats, O&O - Q1 2019 All Hands."
[360] GOOG-NE-06839089 at -94. "Project Bernanke." ("GDN submits two bids into AdX auction…")
[361] During the entire lifetime of Project Bernanke, AdX conducted second-price auctions. This subsection covers the original Project Bernanke, and later subsections cover its variants.
[362] Project Bernanke was launched in 2013. It was in place until Project Global Bernanke was launched in 2015. *See* GOOG-DOJ-28385887 at -93, 94. August 17, 2015. "Beyond Bernanke." ("Bernanke (late 2013)… Global Bernanke [redacted]
[redacted]
[365] To be clear, I am opining that this aspect of Project Bernanke, and collusion among GDN bidders, helps GDN at the expense of publishers. This is not an opinion on Project Bernanke as a whole.
[redacted]
[367] To be clear, I am opining that this aspect of Project Bernanke, and overbidding in a second-price auction, helps publishers at the expense of GDN advertisers – this is not an opinion on Project Bernanke as a whole.

ii. [redacted]

235. Project Bernanke computes the pair of adjustment parameters (α,β) using historical GDN data[368] and auction simulations[369] on a per-publisher basis in a manner that maintains an average take rate of 14% for each billing period. In order to maintain an average take-rate of 14%, Google created for each publisher a “Bernanke pool” that is added to whenever GDN’s take rate exceeds 14% and consumed from whenever GDN’s take rate falls below 14%.

236. An excerpt in Figure 39 from an internal Google document below shows how GDN formulated Project Bernanke to maximize GDN revenue while maintaining a take rate of 15% per publisher.

[368] GOOG-NE-13468541 at -42. “Bernanke experiment analysis.” (“It is important to note that in this entire process, we only use information about the GDN bid and the GDN price paid on queries won by GDN.”)
[369] GOOG-NE-13468541 at -42. “Bernanke experiment analysis.” (“The optimal combination of first bid increase and second bid decrease for each publisher is estimated using AdX auction simulations...”)

Figure 39: An excerpt from a Google slide deck stating that α and β were chosen to maximize GDN revenue while maintaining a take rate of 15%[370]

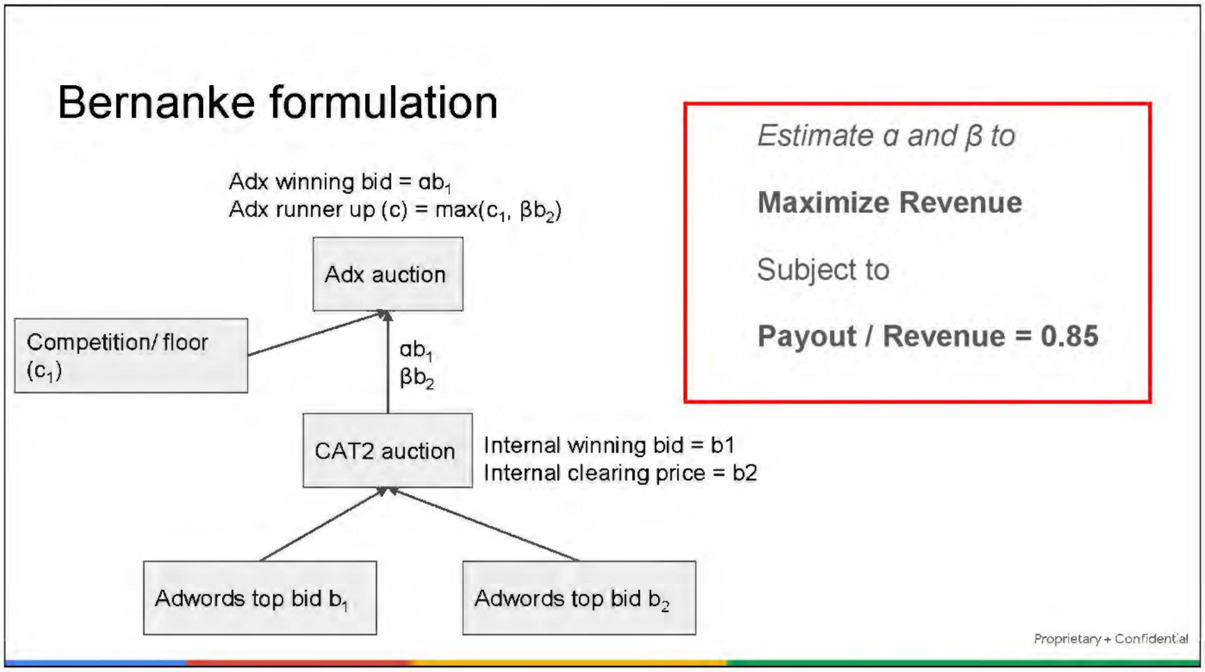

237. Project Bernanke results in an increased win rate and revenue for GDN. I explain why this is the case below as well. Because GDN submits a higher top bid to AdX under Project Bernanke, this means that GDN clears more impressions and therefore has a higher win rate. Moreover, Project Bernanke collects a weakly higher payment from its top advertiser on every single impression. In particular, if GDN would have cleared the impression without Project Bernanke, it collects exactly the same payment. If GDN clears the impression only due to Project Bernanke, it collects non-zero payment (as opposed to zero without Project Bernanke). Because Project Bernanke balances GDN's revenue as 14% of its received payment, GDN also sees increased revenues.

238. The excerpts in Figure 40 from an internal Google document[371] present examples of how Project Bernanke works. In both slides, GDN's top bid into AdX is raised to $1.20, even though the top GDN value is $1.00. This can be interpreted as overbidding, which withdraws money from the Bernanke pool. In both slides, GDN's second bid into AdX is decreased to $0.48 instead of $0.83 (if GDN submitted the true value of the second-highest bidder less a 14% take-rate on

[370] GOOG-NE-11753797 at -37. February 11, 2019. "DVAA Quality, Formats, O&O - Q1 2019 All Hands." (at the time of this document in 2019, GDN take rate was 15%.)
[371] GOOG-NE-06839089 at -98. October 2013. "Project Bernanke - Quantitative Easing on the AdExchange."

$0.96, this would be $0.83). This can be interpreted as collusion. Lowering the second-highest bid, after learning that there is a higher GDN bid, helps the highest GDN bidder. In "Project Bernanke: Example #1", overbidding is necessary to win, as the without it the top GDN bidder would not have cleared the publisher's reserve of $1.10. In this case, GDN must pay AdX $1.10, but can only charge its top advertiser at most $1.00, and therefore this transaction consumes $0.10 from the 'Bernanke pool' to subsidize overbidding. In "Project Bernanke: Example #2", GDN second-prices itself, and the collusion has a positive impact on GDN. Instead of paying AdX $0.83, GDN only pays AdX $0.48 and can pocket the difference into its 'Bernanke pool' to subsidize more overbidding. In the long run, Google structured Project Bernanke so that these additions and withdrawals cancel one another out, allowing GDN to win auctions it otherwise would not have won. Appendix H contains numerical examples of these dynamics.

Figure 40: Excerpts from an internal Google slide deck show that Project Bernanke dynamically adjusts the highest and second highest GDN bids before sending them to the AdX auction[372]

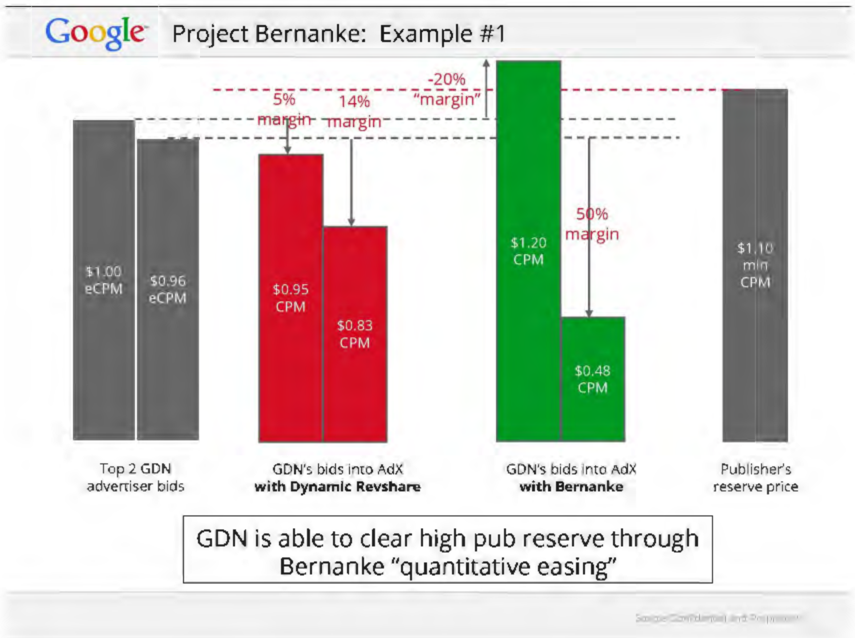

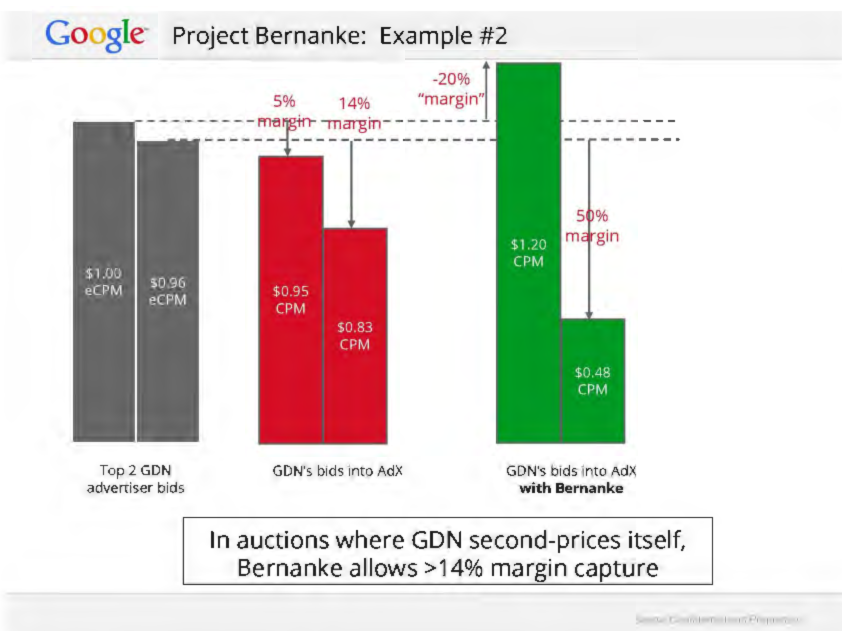

B. Project Global Bernanke

239. Project Bernanke was later expanded to **Project Global Bernanke** in 2015,[373] which instead targeted an average take-rate of 14% across all publishers [redacted]

[redacted][374]

240. Project Global Bernanke maintained a single pool for all publishers across AdX instead of individual per-publisher pools. Project Global Bernanke aimed to keep this single pool roughly empty at the end of each billing period, implying an average take rate of 14% across all of AdX, subject to the aforementioned additional constraints such as floors on the revenue of individual publishers and on the average take rate for individual publishers.

241. Project Global Bernanke also certainly increased GDN's win rate and revenues, for exactly the same reasons as Project Bernanke. I elaborate on why this is the case below.

C. Impact of Projects Bernanke and Global Bernanke on Publishers

242. Under Project Bernanke, individual publisher revenues may increase, may decrease, or may stay the same. The impact on a publisher depends predominantly on whether Project Bernanke results more often in 'poaching' an impression from other ad buying tools or in clearing impressions that otherwise would have gone unsold because no bids surpassed the reservation price.

243. More specifically, with per-publisher pools, when Project Bernanke's overbidding causes GDN to win an impression it otherwise would not have, the impact of these impressions on a publisher depends on the comparison of (a) the payment AdX would have received without Project Bernanke, and (b) the highest pre-Bernanke GDN bid for this impression. If (a) exceeds (b) on average (*i.e.,* if Project Bernanke largely results in GDN 'poaching' impressions that would otherwise have been won by other ad buying tools), then the publisher faces a revenue loss. If (b) exceeds (a) on average (*i.e.,* if Project Bernanke largely results in GDN winning impressions that otherwise would have gone unsold), then the publisher revenue increases. This is because each publisher loses revenue when the Bernanke pool grows (and by exactly the amount the pool grows), and gains revenue when the Bernanke pool shrinks (and by exactly the amount the pool shrinks), plus the difference between (b) and (a).[375] Because Project Bernanke constrains each publisher to have a neutral pool on average, the cumulative average increase/decrease to the

[372] GOOG-NE-06839089 at -98, 99. October 2013. "Project Bernanke - Quantitative Easing on the AdExchange."
[373] Also called Project Bell v1. *See* GOOG-AT-MDL-006218257 at -63. December 16, 2022. "Case AT.40670 - Google - Adtech and Data-related practices." ("Project Bell was also known as Global Bernanke.")
[374] GOOG-DOJ-AT-02471194 at -4. July 26, 2015. [redacted]
[375] See Tables 1, 2 and 3 in Appendix H.5.

publisher's revenue from these impressions comes exactly via the difference between (b) and (a).[376] Therefore, the change in revenue earned by publishers through impressions that GDN wins through AdX is determined exactly via the difference between (b) and (a).

1) Project Global Bernanke led to revenue reduction in some publishers

244. Under Project Global Bernanke, individual publisher revenues may increase, may decrease, or may stay the same. The reasoning is the same as the paragraph above, plus one additional complexity due to the global pool. Specifically, Project Global Bernanke creates a further complication with its single global pool under which some publishers may incur losses simply because they contribute more to the Global Bernanke pool than the payouts they receive from it. Stated differently, some publishers may see decreased revenues by Project Global Bernanke because their impressions may have contributed to the Bernanke pool on average, whereas other publishers may have consumed from the pool on average. As a result, it is still the case that publishers for whom (a) is larger than (b) tend to see decreased revenues under Project Global Bernanke, and those for whom (b) is larger than (a) tend to see increased revenues. As compared to the original Project Bernanke, under Project Global Bernanke it is also the case that publishers who tend to contribute to the Global Bernanke pool (*i.e.*, AdX often provides both the highest and second highest bid for such publishers' impressions) more than they consume (*i.e.*, AdX often is not the highest bid for such publishers' impressions) may also see decreased revenues under Project Global Bernanke compared to Project Bernanke, whereas those publishers that consume more than they contribute tend to see increased revenues from Project Global Bernanke as compared to Project Bernanke.

245. In sum, the above arguments imply that:

a. When comparing the revenues a publisher sees under Project Bernanke versus no Bernanke, this predominantly comes down to comparing the average across all impressions cleared due to Project Bernanke of (a) payment AdX would have received without Project Bernanke and (b) the highest pre-Bernanke GDN bid.

[376] In addition, Projects Bernanke and Global Bernanke can also cause a non-GDN advertiser to pay more for an impression that they still win under Projects Bernanke and Global Bernanke. For example, imagine that a non-GDN advertiser bids $10 on an impression, and the second-highest bid is $4 from GDN (without Project Bernanke). With Project Bernanke, the second-highest bid might be increased to $8. This is not high enough to claim the impression, but causes the non-GDN advertiser to pay more, and the publisher to claim an extra $4 in revenue.

Publisher revenues predominantly decrease if (a) is higher than (b), and predominantly increase if (b) is higher than (a).[377]

b. When comparing the revenues a publisher sees under Project Global Bernanke versus Project Bernanke, this is predominantly determined by comparing whether the publisher contributes more than they consume from the Global Bernanke pool.[378] Publisher revenues predominantly decrease if they contribute more than they consume, and predominantly increase if they consume more than they contribute.

c. When comparing the revenues a publisher sees under Project Global Bernanke versus no Bernanke, a publisher for whom (a) is greater than (b) and who contributes more than they consume certainly sees predominantly decreased revenues, and a publisher for whom (b) is greater than (a) and who consumes more than they contribute certainly sees predominantly increased revenues. A publisher for whom (a) is greater than (b) but who consumes more than they contribute, or for whom (b) is greater than (a) but contributes more than they consume, would require a quantitative comparison to ultimately determine their change in revenue from no Bernanke to Project Global Bernanke.

246. Google internal documents show these disparate effects of Project Global Bernanke on publishers as compared to Project Bernanke. [redacted]

[377] As previously noted, another relevant factor is the increased revenue due to Project Bernanke increasing the payment of non-GDN winners. Put another way, the revenue paid by GDN to AdX increases or decreases entirely based on a comparison of (a) to (b), but the payout of non-GDN advertisers to AdX could increase.
[378] An additional effect is that Project Global Bernanke formulates a global optimization rather than per-publisher optimizations, so the set of impressions impacted by Project Global Bernanke will change, and the precise multipliers will also change, as compared to Project Bernanke. For example, due to the flexibility of global versus per-publisher optimization, more impressions were likely impacted by Project Global Bernanke than Project Bernanke. I say 'predominantly' to indicate that this is the novel form of impact when comparing Project Bernanke to Project Global Bernanke, but to not preclude other forms of impact.
[379] GOOG-DOJ-AT-02471194. July 26, 2015. "Global Bernanke."

2) Projects Bernanke and Global Bernanke can lead to a reduction in ad quality as well as revenue per mille for publishers

247. If GDN advertisers tend to display lower quality ads, Projects Bernanke and Global Bernanke would lead to lower quality ads displayed for publishers. Because GDN wins more impressions under Projects Bernanke and Global Bernanke, if GDN advertisers tend to display lower quality ads than non-GDN advertisers, then this program would cause publishers to receive lower quality ads because it causes GDN advertisers to win impressions more frequently.

248. Even publishers who saw a revenue increase under Projects Bernanke and Global Bernanke can see their revenue per mille ("RPM") decreased. As shown by an internal Google presentation,[380] Project Bernanke led to a decrease in RPM across publishers. This decrease would be because a publisher can benefit from Projects Bernanke and Global Bernanke on a per-auction basis only when an impression is sold that would otherwise have gone unsold because no bids exceeded the reservation price. Hence, a publisher who is concerned about RPM might still prefer the outcomes under no Bernanke as compared to Projects Bernanke and Global Bernanke, even if their overall revenue does not decrease.

3) Publishers would have raised their reserve prices to maximize their revenue had they known about Projects Bernanke and Global Bernanke

249. Google concealed vital information from publishers by concealing Projects Bernanke and Global Bernanke. In particular, there are two tests a publisher could do to determine how Projects Bernanke and Global Bernanke predominantly impact their revenues compared to no Bernanke.

a. As noted earlier, publisher revenue might increase, decrease, or be neutral under Projects Bernanke and Global Bernanke. Under Project Bernanke, a publisher benefits on impressions purchased by GDN only when Project Bernanke largely results in GDN winning impressions that otherwise would have gone unsold rather than in GDN winning inventory where a non-GDN buyer previously set the clearing price. A publisher with a low rate of unsold inventory, or a high volume of competitive non-GDN bids, would likely see decreased revenues under Project Bernanke and prefer their outcomes without Project Bernanke.

b. Additionally, under Project Global Bernanke, publishers who tend to have GDN provide both the highest and second highest bid more often tend to contribute to

[380] GOOG-DOJ-28386151 at -61. December 10, 2013. "Project Bernanke - Quantitative Easing on the AdExchange."

the Global Bernanke pool rather than consuming from it.[381] Such publishers would prefer their outcomes without Project Global Bernanke.

250. Moreover, all publishers likely would have changed their behavior if they knew about Projects Bernanke and Global Bernanke by raising their reserve prices.[382] For a publisher to see predominantly increased revenues from Project Bernanke, they must tend to interact with it primarily through GDN purchasing unsold inventory rather than GDN stealing impressions from another ad buying tools, and for Project Global Bernanke, they must further interact primarily through GDN purchasing unsold inventory rather than GDN submitting the top two bids. A publisher who raises their reserves causes (a) more inventory to go unsold, increasing the positive interactions with Project Bernanke, (b) non-Google ad buying tools to win less often (because they may fail to clear the reserve), and (c) GDN to be less likely to submit top two bids (because the second price is less likely to clear the reserve). Hence, raising reserves would cause a publisher to primarily have revenue-boosting rather than revenue-draining interactions with Project Bernanke. Under Project Global Bernanke, raising reserves would further cause a publisher to primarily have pool-consuming interactions (which increase publisher revenue) rather than pool-contributing interactions (which decrease publisher revenue) with Project Global Bernanke.

D. Impact of Projects Bernanke and Global Bernanke on Ad Buying Tools

1) Under Projects Bernanke and Global Bernanke, GDN increased its revenue at the expense of non-Google ad buying tools

251. Under Project Bernanke and Global Bernanke, GDN increases its revenues and win rate. GDN maintains an average 14% take rate of all revenue paid by its advertisers in AdX auctions. Under Projects Bernanke and Global Bernanke, its advertisers make weakly higher payments on every single instance (which therefore increases GDN's win rate) because:

a. If GDN would have won without Projects Bernanke or Global Bernanke, the winning GDN advertiser pays the same. The only difference under Project Bernanke is that some portion of this payment now goes to the Bernanke pool.

[381] Whereas publishers who tend to sell impressions that GDN would not win without Project Global Bernanke (either because a non-GDN advertiser wins, or because the impression would go unsold) would consume more than they contribute.
[382] Importantly, Reserve Price Optimization, which I analyze in Section IX, does not apply to GDN. GOOG-TEX-00831090 at -1. April 17, 2017. "DRX 2.0 Quality." (see table.)

b. If GDN would not have won without Projects Bernanke or Global Bernanke, the winning GDN advertiser pays their bid which is certainly larger than $0 they would have paid if there were no Project Bernanke or Global Bernanke.

252. Google internal documents provides measurements of GDN benefits created by Projects Bernanke and Global Bernanke. [redacted]

253. Under Projects Bernanke and Global Bernanke, the win rate of non-GDN advertisers and non-GDN ad buying tools on AdX would decrease. This follows immediately because GDN submits a weakly higher bid for every AdX impression under Projects Bernanke and Global Bernanke as compared to no Bernanke. In addition, non-GDN advertisers would pay more for any impressions they still win under Projects Bernanke and Global Bernanke, for the same reason, but their total aggregate payment could still decrease due to winning fewer impressions.[386]

254. Google internal documents confirm this analysis as well. [redacted]

---

[383] GOOG-NE-03872763 at -81. "Discussion on improving AdX and AdSense backfill."
[384] GOOG-DOJ-28385887 at -95. August 17, 2015. "Beyond Bernanke."
[385] GOOG-DOJ-28386151 at -60. December 10, 2013. "Project Bernanke - Quantitative Easing on the AdExchange."
[386] GDN impacts non-GDN advertisers and non-ad buying tools only by determining their minimum bid to win. Only GDN's highest submitted bid impacts others' minimum bid to win, and others' minimum bid to win increases with GDN's highest submitted bid. Once non-GDN advertisers have a higher minimum bid to win, they're immediately less likely to win, and also will pay more if they still win.
[387] GOOG-DOJ-AT-02513569 at -73. "gTrade Team Background."
[388] GOOG-DOJ-28386151 at -69. December 10, 2013. "Project Bernanke - Quantitative Easing on the AdExchange."

E. Impact of Projects Bernanke and Global Bernanke on Advertisers

1) Projects Bernanke and Global Bernanke did not benefit GDN advertisers, but decreased win rates for advertisers using non-Google ad buying tools

255. Under Projects Bernanke and Global Bernanke compared to no Bernanke, I would expect GDN advertisers to achieve identical payoff, defined as the difference between their value for impressions won and payments made.[389] In particular, the only change that GDN advertisers see under Project Bernanke and Global Bernanke is that they now sometimes win impressions and pay their value (for a net payoff of $0) whereas they would have lost without Project Bernanke (which yields a payoff of $0).[390] More specifically,

a. In auctions where Projects Bernanke and Global Bernanke decrease the second highest GDN bid to replenish the Bernanke pool, the advertiser wins and pays the second highest bid in full (even though the amount GDN pays to AdX decreases). Had there been no manipulation, the advertiser would have still won and paid the second highest bid again. As a result, these auctions have no impact on advertisers.

b. In auctions where GDN wins only due to Projects Bernanke or Global Bernanke increasing its highest bid, the GDN advertiser wins and pays their bid, leading to a payoff of $0.[391] Had there been no manipulation, they would have lost the impression, either to some non-GDN advertiser or because their bid was below the reserve, and paid nothing. This would lead to a payoff of $0 as well. As a result, these auctions do not impact GDN advertisers' payoffs either.

Hence, overall, Projects Bernanke and Global Bernanke were neutral towards GDN advertisers' payoffs.[392]

---

[389] This subsection assumes that advertisers bid their true value, which is a dominant strategy in a second-price auction with reserve (see Section II). Because Projects Bernanke and Global Bernanke were never revealed to advertisers, this is reasonable behavior to expect. The conclusions I draw still hold approximately if instead some advertisers slightly shaded their bids in response to Bernanke or Global Bernanke.

[390] If an advertiser is interested in Return on Investment (ROI) in addition to payoff, that advertiser would prefer no Bernanke to Project Bernanke or Global Bernanke. Specifically, winning an impression valued at v for price v is neutral towards payoff, but harms average ROI (because the ROI on this impression is as low as possible while paying at most the value).

[391] Again, this assumes that GDN advertisers report their true value to GDN, which is the dominant strategy in a second-price auction (and AdX was described as a second-price auction).

[392] There are two caveats to this. First, advertisers who considered shading the bid they entered to GDN, even though AdX was described as a (truthful) second-price auction, may have seen increased payoff. Second, advertisers who value ROI in addition to payoff could still have preferred outcomes without Projects Bernanke or Global Bernanke, even with a neutral effect on payoff.

256. Under Projects Bernanke and Global Bernanke, the win rate of non-GDN advertisers on AdX would decrease. This is because the GDN advertisers are still winning every impression that they would have won without Projects Bernanke and Global Bernanke, but they are also winning additional impressions. Some of these impressions previously would have been won by non-GDN advertisers, so these advertisers face a lower win rate.[393]

257. Google internal documentation shows that non-GDN advertisers saw a decline in their win rates. [redacted] This decrease in spending can only be explained by a decrease in their win rates.[396]

2) Advertisers would have shaded their bids to maximize their payoff had they known about Projects Bernanke and Global Bernanke

258. Google concealed vital information from advertisers by concealing Projects Bernanke and Global Bernanke. Provided that neither Project Bernanke nor Project Global Bernanke were disclosed to advertisers, they would naturally believe they were still participating in a truthful second-price auction and bid their true value as a result. If advertisers knew they were participating in a non-truthful auction, they would have instead considered shading their bids. Knowing the auction format is vital information to advertisers aiming to optimize their payoff. In particular, Projects Bernanke and Global Bernanke are dirty second-price auctions.[397] Specifically, if c denotes the minimum bid to win for GDN on AdX, then from the perspective of a GDN advertiser, Projects Bernanke and Global Bernanke are both dirty second-price auctions with soft floor c and hard floor c/α. That is, as long as the highest GDN bidder exceeds c/α, they will win, because their bid will be increased by α to exceed c. If their bid further exceeds c, the auction turns into a regular second-price auction. If their bid falls between c/α and c, they will pay their bid and be subsidized by Bernanke pool for the remaining amount. Therefore, this is a dirty second-price auction.

[393] Some of these impressions could have been previously unsold.
[394] GOOG-DOJ-28385887 at -95. August 17, 2015. “Beyond Bernanke.”
[395] GOOG-DOJ-28386151 at -67. December 10, 2013. “Project Bernanke - Quantitative Easing on the AdExchange.”
[396] This is because, as noted previously, the price a non-GDN advertiser pays if they still win under Projects Bernanke and Global Bernanke can only increase from what they would pay with no Bernanke.
[397] See Section VII for more information on dirty second-price auctions.

259. Moreover, it is possible that some advertisers preferred outcomes under no Bernanke as compared to Projects Bernanke or Global Bernanke. Advertisers might be concerned about ROI in addition to payoff. Projects Bernanke and Global Bernanke result in lower ROI for GDN advertisers, because GDN advertisers spend more but without generating positive returns (by purchasing impressions at a price exactly equal to their willingness to pay). Therefore, an advertiser who prefers a higher ROI option among two with equal payoffs would prefer no Bernanke to Projects Bernanke or Global Bernanke.

260. In addition, advertisers may choose to use GDN simply to avoid the negative impacts caused by Projects Bernanke and Global Bernanke to non-GDN advertisers, and not because of any benefits they receive from Project Bernanke or Global Bernanke. In particular, I have previously described that Projects Bernanke and Global Bernanke are neutral towards GDN advertisers' payoffs, and only harm ROI. On the other hand, I have also previously described that Projects Bernanke and Global Bernanke decrease non-GDN advertisers' payoffs (both by decreasing their win rate, and by increasing their payment per-impression). Therefore, had they known about the conducts, an advertiser might prefer GDN to non-GDN simply to avoid the negative impacts caused to non-GDN advertisers by Projects Bernanke and Global Bernanke.

261. Under Projects Bernanke and Global Bernanke, market efficiency may increase, may decrease, or may be stable. Overall efficiency is maximized when the impression is awarded to the highest value bidder. Projects Bernanke and Global Bernanke affects this in two ways. First, if the impression would have been sold to a non-GDN advertiser but is now awarded to a lower value GDN advertiser, overall efficiency decreases because the impression is taken away from a higher-value bidder and transferred to a lower-value bidder. Second, if the impression would have been unsold but is now awarded to a GDN advertiser, overall efficiency increases because the impression is now sold when it otherwise would have gone unsold. In particular, the impact on the overall efficiency of the market is entirely determined by comparing two quantities, on average, when overbidding causes GDN to win an impression it otherwise would not have: (a) the value of the winning advertiser without Project Bernanke ($0 if the impression would have been unsold), and (b) the highest GDN bid for this impression. If (a) exceeds (b) on average, then the overall market efficiency decreases. If (b) exceeds (a) on average, then the overall market is more efficient. Note, in particular, that because the value of the winning advertiser without Project Bernanke is always higher than the payment AdX would have received without Project Bernanke, it is possible for the overall market efficiency to decrease even if publisher revenue on average increases.

F. Impact of Projects Bernanke and Global Bernanke on Exchanges

262. When combined with (Enhanced) Dynamic Allocation, Projects Bernanke and Global Bernanke enabled AdX to have a higher win rate, which would cause other exchanges to have a lower win rate.[398] Under (Enhanced) Dynamic Allocation, AdX won every impression for which it exceeds its reserve. Projects Bernanke and Global Bernanke inflate the highest bid to AdX, making AdX more likely to clear its reserve. Therefore, AdX is more likely to win. This holds both when other exchanges participate in the waterfall process (because fewer impressions will even continue down the waterfall if AdX wins them early), and when other exchanges participate in header bidding (because AdX is now more likely to exceed the highest header bidding bid, which acts as AdX's reserve). Since AdX is more likely to win, this directly implies that other exchanges are more likely to lose, due to the fixed number of impressions.[399]

263. When AdX participates in a simultaneous auction with other exchanges, Projects Bernanke and Global Bernanke could cause AdX to win more or fewer impressions. When AdX uses a second-price auction to participate in an auction with other exchanges, such as in Exchange Bidding, it matters not only whether AdX clears its reserve, but also at what price it clears. A higher clearing price would cause AdX to win more often, and a lower clearing price would cause AdX to win less often. Under Projects Bernanke and Global Bernanke, the overbidding aspect causes AdX to have a higher clearing price, and therefore would cause AdX to win more often. On the other hand, the collusion aspect causes AdX to have a lower clearing price, and therefore would cause AdX to win less often.[400] Because impacts are possible in both directions, AdX would sometimes have a higher clearing price and sometimes have a lower clearing price.[401]

[398] This conclusion assumes that publishers did not significantly inflate AdX's price floor due to Projects Bernanke and Global Bernanke. This seems plausible, due to (a) Projects Bernanke and Global Bernanke were never disclosed and (b) for publishers who let (Enhanced) Dynamic Allocation set AdX's price floor without an additional 'boost', AdX's price floor would be the maximum Value/temporary CPM of other exchanges/guaranteed line items, which are not impacted by GDN's bids on AdX.
[399] In particular, exchanges that participate via header bidding are more likely to have their bids superseded by AdX, and exchanges that participate via the waterfall are more likely to see impressions taken by AdX before having an opportunity to solicit bids.
[400] Note that the collusion aspects cannot cause AdX to fail to meet its reserve. But conditioned on AdX meeting its reserve, it could cause AdX to have a lower clearing price, which is exactly how GDN benefits.
[401] Note that under (Enhanced) Dynamic Allocation, the magnitude of AdX's clearing price does not impact whether it wins the impression. It only matters whether AdX beats its reserve or not.

G. Project First-Price Bernanke

264. AdX eventually switched to a first-price auction,[402] which renders the particular collusion mechanics of old Project Bernanke obsolete. This is because the first-price auction is pay-your-bid, hence dropping GDN's second highest bid does not impact the auction at all. The general framework of colluding and overbidding still apply, but the precise mechanics differ. I provide more details on this in Appendix H.

265. As I previously discussed, first-price auctions are not truthful. In fact, for any bid b less than the bidder's value *v*, it is always better to bid b instead of v.[403] However, an economic approach called the "Revelation Principle"[404] allows an intermediary to make a first-price auction truthful for participants. Intuitively, it works in the following way: The intermediary first comes up with a device that takes in a bidder's value as an input and calculates the optimal bid. The intermediary then tells the bidders to report their true values and assures them that if they win, they will be charged their minimum bid to win. But the intermediary does not submit the true values of the bidders to the auction on their behalf, and instead submits bids calculated by the optimization device. The auctioneer then executes a first-price auction with those bids. From the perspective of the advertisers, this is a truthful auction, because they always pay their minimum bid to win. But potentially there is a mismatch in payments since the minimum bid to win (what is charged to the winning bidder) might differ from what is calculated as the optimal bid from the highest value submitted by the bidders (what is paid to the auctioneer). If the device is excellent at bid optimization, these will perfectly balance out on average. If not, there can be a benefit or loss to either party. For the rest of this analysis, I assume that the bid optimizer is excellent.[405]

266. First-Price Project Bernanke has three components: (a) a bid optimizer for GDN users that makes their participation in AdX's first-price auction truthful, (b) collusion among GDN bidders, which increases GDN's payoff at the expense of publishers' revenue, (c) overbidding, which lowers GDN's and increases publishers' revenue. In comparison to Projects Bernanke and Global

[402] *See* Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20240122142404/https:/blog.google/products/admanager/update-first-price-auctions-google-ad-manager/
[403] To see this, observe that bidding v guarantees a payoff of 0, no matter what since it either leads to losing the item, or winning the item and paying the bid. Bidding b that is less than v instead guarantees a payoff no worse than 0 since it either leads to losing the item or winning the item and paying less than the value.
[404] *See* Roger B. Myerson. "Incentive Compatibility and the Bargaining Problem." *Econometrica* vol. 47, no. 1. 1979. pg. 61–73.
[405] Shortcomings of the bid-optimizer are certainly relevant for thinking through the impacts of First-Price Project Bernanke, but it is not relevant to the conclusions I draw based on collusion and overbidding alone.

Bernanke, (b) and (c) are conceptually similar but implemented via different mechanics (due to the different mechanics between first- and second-price auctions).

267. First-Price Project Bernanke carries the same motivation as Project Bernanke. GDN bidders could collude in AdX's auction to increase GDN's revenue at the expense of the publisher, but this hurts publishers' revenues. First-Price Project Bernanke again observes that overbidding has the opposite effect of lowering GDN's payoff but helping publishers' revenues, although it causes collateral damage to non-GDN advertisers. Additionally, there is also an added complication due to intermediating AdX's first-price auction to make it truthful.

268. Under First-Price Project Bernanke, GDN manipulated advertisers' bids before sending them to AdX in the following manner:[406, 407, 408]

a. [redacted]

[redacted]

[redacted]

[redacted]

[406] First Price Project Bernanke was launched in 2019. *See* GOOG-AT-MDL-008842383 at -88. August 5, 2023. "Declaration of Nirmal Jayaram." ("Google updated the Bernanke algorithms in 2019 to be compatible with the Unified First Price Auction. The updated version of Bernanke was sometimes referred to within Google as 'Alchemist.'")
[407] GOOG-DOJ-AT-02224828. March 2019. "The Alchemist." [redacted] My best guess is that the core ideas have not changed.
[408] This was also called "The Alchemist." [redacted]

- [redacted]
- [redacted]

269. To illustrate how First-Price Project Bernanke works, imagine that there are ten total bidders, and each is believed to have a valuation uniformly in the range [0,10]. If all players use an ad buying tool with a 15% take rate, the best course of action is for each ad buying tool to submit a bid of 0.9*0.85*v to AdX when their advertiser's value is v.[412] If an ad buying tool wanted to modify the auction process to make it truthful without also implementing First-Price Project Bernanke, they could tell their advertisers that when an advertiser submits a bid of b to the tool, the tool submits a bid of 0.9*0.85*b on their behalf to AdX. If the tool wins, and the highest other bid is H, then the tool pays 0.9*0.85*b to AdX, and charges the winning advertiser H/(0.9*0.85) which corresponds to the minimum bid the advertiser could have submitted to their ad buying tool and still had the winning bid. While a single auction is unlikely to result in a take rate of exactly 15%, this will average out to a take rate of 15% over time.

270. First-Price Project Bernanke additionally implements collusion among its advertisers. Say for example that nine of the advertisers use GDN. [redacted]

[412] This follows as it is a Bayes-Nash equilibrium for each advertiser to bid 0.9*v, and their behavior is not impacted by what take rate the ad buying tools use, provided they all use the same take rate.

[redacted]

a. Collusive aspect. In the above example, without any collusion, each ad buying tool would optimally submit a bid of 0.9*0.85*v on behalf of an advertiser with value v. This is because that advertiser has *uncertainty surrounding the other values of all nine other bidders* and can't shade their bid by too much and still hope to win. If instead, nine of the advertisers get together through GDN and determine who is the highest bidder, now there is only *uncertainty surrounding the other value of just one other bidder*, and it is safe to bid-shade more aggressively. In both a first- and second-price auction, the first step to successful collusion among a team of bidders is to figure out who is the highest bidder. In a second-price auction, the remaining bidders simply drop out in order to lower the winner's eventual payment, but the winning bidder need not adjust their bid as the second-price auction is truthful. In a first-price auction, dropping out alone does not help the winning bidder (because they will still pay their bid if they win, and they will still win if and only if their bid exceeds the highest non-colluding bid). But now that the highest colluding bidder *knows that they face less competition*, they can shade their bid more aggressively to get higher payoff. In the example, this manifests with an optimal bid-shade of 50% versus 90%.[413]

b. Overbidding aspect. In the above example, overbidding represents that the ad buying tool submits the optimal bid for a value of α*b, instead of the optimal bid for b. In this example, to implement overbidding without collusion, the ad buying tool would submit a bid of 0.9*0.85*α*b. With collusion and overbidding, the ad buying tool submits a bid of 0.5*0.85*α*b. This aspect is quite similar to overbidding in a second-price auction, after accounting for the role of the ad buying tool in implementing the Revelation Principle.

[413] Here is a simpler, but less realistic example. Imagine instead that all ten bidders collude in a first-price auction. Once they determine the highest bidder amongst themselves, that bidder should optimally submit a bid of just a penny, and the others will drop out. Note that if the highest bidder submits their original optimal bid of 0.9*0.85*v, then they do not benefit simply because the colluders drop out – the highest bidder must further use this information to re-optimize their bid (which in this case is just a penny, as there is no remaining competition). This reasoning qualitatively extends to the chosen example, and indeed all instances, although the required mathematics is more involved.

271. When AdX participates in a simultaneous auction with other exchanges, First-Price Project Bernanke could cause AdX to win more or fewer impressions. When AdX participates in an auction with other exchanges, such as in Exchange Bidding, it matters not only whether AdX clears its reserve, but also at what price it clears. A higher clearing price would cause AdX to win more often, and a lower clearing price would cause AdX to win less often. In a first-price auction, the clearing price is the highest bid, and therefore AdX's clearing price would increase or decrease based on whether Project First-Price Bernanke causes GDN's highest submitted bid to increase or decrease. The overbidding aspect causes GDN to submit a higher bid. On the other hand, the collusion aspect causes GDN to submit a lower bid.[414] Because impacts are possible in both directions, AdX would sometimes have a higher clearing price and sometimes have a lower clearing price.

IX. CONDUCT ANALYSIS: RESERVE PRICE OPTIMIZATION

272. In this section, I provide an analysis of Google's Reserve Price Optimization (RPO)[415] conduct. I demonstrate that RPO leads to higher revenue for Google's ad exchange AdX and explain the mechanisms through which it leads to lower payoff to advertisers and could lead to lower revenue for some publishers. Furthermore, I outline how it impacts publisher and advertiser behavior. The negative effects of RPO to advertiser payoff, and possibly some publishers' revenues, at least partially stem from Google's concealment of the conduct.

273. In particular, it is my opinion that Google concealed information that is material to both publishers and advertisers during the period RPO was concealed. It is also my opinion that even after RPO was revealed, publishers might set suboptimal reserves on any impression for which RPO is a possibility.[416]

274. Under RPO, AdX used data available to them (prior to seeing live bids) to calculate per-buyer reserve prices[417] that it believed would optimize AdX's revenue.[418] AdX then used these

[414] Note that the collusion aspect cannot cause GDN to fail to submit a bid above AdX's reserve. But conditioned on meeting AdX's reserve, GDN could submit either a higher or lower bid.
[415] There seems to be other programs that are called RPO previously. However, they substantially differ from the conduct I am discussing here. The main difference between those programs and this conduct is that this conduct sets per-buyer reserve prices.
[416] That is, unless a publisher knows whether RPO activates on a particular impression, and if so what reserve RPO would set, I would expect publishers to lack sufficient information to set a profit-maximizing reserve.
[417] GOOG-NE-13204729 at -36. August 17, 2015. "AdX Dynamic Price." ("Buyer Specific Reserve Prices. Different buyers may get different reserve prices.")
[418] GOOG-NE-06151351 at -52. November 12, 2015. Email thread, "Subject: [Monetization-pm] Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers." ("[RPO] generate[s] a histogram of historical bids and transaction prices […] pick[s] a reserve price that maximizes predicted revenue.")

reserves in its own auction instead of the reserves set by the publisher, although this reserve was always at least as large as the reserve set by the publisher.[419] The program was launched in phases between April and October 2015.[420] Initially, Google did not announce this program to its customers.[421] Later, Google announced the program to its customers under the name "optimized pricing" on May 12th, 2016, more a year after its initial rollout.[422] Publishers were not allowed to opt out of the program.[423] The program was deprecated in 2019 with the switch of AdX to the first-price auction format.[424]

275. Internal Google documents suggests that RPO relies on an algorithmic optimization that "set[s] optimized reserve prices in AdX auction[s]" to "increase the revenue for publishers" via "model[ing] effect of various reserve prices" and then "pick[ing] the best one."[425] Importantly, this tool aims to set the reserve price just below what the highest bidder is willing to pay,[426] by coming up with an empirical estimate of this willingness to pay, which was assumed to be equal their bid due to the truthfulness of the second-price auction.[427] An internal Google document states that the goal of RPO was to "select a reserve price as close to the anticipated first price as possible in order to trade buyer for seller surplus."[428] If AdX has sufficient data to form an accurate prediction of the maximum advertiser value v, the optimal reserve price to set is exactly v.

Different Google internal documents outline different strategies AdX used to employ the data they have to best estimate the RPO reserve prices. Which data was used and how data was processed are not relevant to the conclusions I provide below.

[419] GOOG-NE-03640022 at -2. "AdX Managed Reserves." ("Currently RPO can only raise reserve prices".) Notice that optimizing publisher revenue and AdX revenue are equivalent since AdX revenue corresponds to 20% of the publisher revenue.

[420] *See, e.g.,* GOOG-NE-06151351 at -52. November 12, 2015. Email thread, "Subject: [Monetization-pm] Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers." ("Between April and October we launched and improved new systems to dynamically set auction reserve prices for AdX sellers.")

[421] GOOG-NE-09485306 at -432. December 18, 2017. "OLD – New Ad Manager Indirect Notes." ("We are not commercializing this externally for now.")

[422] *See* Jonathan Bellack. "Smarter optimizations to support a healthier programmatic market" (May 12, 2016). Accessed on May 31, 2024.
https://web.archive.org/web/20200929015943/https:/blog.google/products/admanager/smarter-optimizations-to-suppor/

[423] GOOG-NE-06842715 at -18. May 10, 2016. "AdX Auction Optimizations." ("No opt-out possible.")

[424] GOOG-AT-MDL-000987708 at -8. April 9, 2021. "PM Perspective on 1P RPO." ("When we transitioned to a 1st price auction and launched unified pricing rules in September 2019, we had to turn off 2P RPO since it was designed to work in a 2nd price auction (duh).")

[425] GOOG-NE-13204729 at -30. August 17, 2015. "AdX Dynamic Price."

[426] When there is sufficient data to predict the highest bidder's willingness to pay exactly, the goal is indeed to set the reserve price just below this. Often there is insufficient data to predict the highest bidder's willingness to pay exactly, and in these cases Reserve Price Optimization instead aims to set the optimal reserve given the information it has.

[427] GOOG-NE-13204729 at -34. August 17, 2015. "AdX Dynamic Price." (the slide titled "How to Guess the Top Bid" explains how Google thought about estimating the highest bid.)

[428] GOOG-NE-03640022 at -2. "AdX Managed Reserves."

276. To illustrate how RPO works, imagine an impression arrives for a user who is over the age of 25, likes toys, and lives in Plano, TX. This coarse targeting data is sufficient for the publisher to estimate that such users on average fetch $5, although certainly there is high variance depending on more fine-grained third-party cookie data. AdX is called on this impression and learns from third-party cookie data that this user further has no children and is working as an unpaid intern. Under RPO, AdX would guess from the third-party cookie data that this impression would not sell for more than $1, so it leaves the publisher's reserve of $5 intact. When AdX solicits bids, the highest bid is $2, and the impression does not clear. Another impression arrives for a user who is over the age of 25, likes toys, and lives in Plano, TX. Now imagine the same setting, but the third-party cookie data shows this user has two children, buys toys for them frequently, and lives in a wealthy neighborhood. Under RPO, AdX would guess from the third-party cookie data that this impression would fetch $50, so it increases the reserve to $50. In the AdX auction, perhaps the two highest bids are $60 and $55. In this case, with or without RPO, AdX's auction clears at $55, netting $11 for AdX and $44 for the publisher. Perhaps instead the two highest bids are $60 and $45. In this case, with RPO, AdX's auction clears at $50, netting $10 for AdX and $40 for the publisher. Without RPO, AdX's auction would have cleared at $45, netting $9 for AdX and $36 for the publisher. In this case, RPO increases AdX's and publisher's revenue, but the advertiser suffers from a higher payment. Note that in all cases, the advertiser suffers under RPO due to facing a higher reserve. In these particular examples, both publisher and AdX profit modestly under RPO, although it is possible for both to profit significantly (at the expense of the advertiser's payoff, in case that the advertiser might have paid way below their value without RPO), or to have significantly decreased profit (at significant cost to the advertiser as well, if RPO accidentally sets too high of a reserve and the impression no longer clears).

A. Impact of Reserve Price Optimization on Publishers

277. Under general circumstances, because RPO can only increase the publisher reserve, it could increase publishers' payoff by (a) generating greater expected revenue from sales if the publisher believes that Google is better at optimizing revenues than themselves, perhaps due to better data and more sophisticated algorithms and (b) causing AdX to be less likely to win impressions, which might improve the publisher's payoff if AdX ads are typically of low quality.[429]

[429] A higher reserve for an auction leads to a lower probability of the item being sold in either a first- or second-price auction.

278. Google internal documents reveal that RPO increased publisher revenues. A launch email states that the program generated an annual revenue increase [redacted] for publishers.[430] Increased revenues notwithstanding, it is still my opinion that concealing RPO from publishers conceals material information from publishers, for the reasons outlined below.

1) RPO can prevent publishers from optimizing their revenue

279. With RPO, Google concealed material information from publishers. As noted in Section II, reserve prices are material to a publisher's revenue. As one example, if Google is good at optimizing reserves via RPO, a publisher may wish to lower the reserve it sets on AdX in order to give AdX greater flexibility in optimizing its reserve, which would lead to greater revenues for both AdX and the publisher. But, via concealing RPO, Google prevented the publishers from effectively optimizing revenue. In fact, even after RPO is disclosed, publishers would still face challenges setting optimal reserves under RPO. For example, if a publisher wishes to lower the reserve it sets on AdX to give RPO more flexibility, they would want to know exactly on which auctions RPO is active. If a publisher prefers to trust their own optimization over Google's, they would further want to know not only whether RPO is active, but also exactly what reserve RPO is setting (so they can set what they consider to be the optimal reserve exceeding RPO's).

280. However, there are also reasons why some publishers might prefer outcomes without RPO than with RPO. For example, a publisher whose goal is to maximize the number of sold ads (as opposed to maximizing their revenue) would prefer outcomes without RPO as opposed to with RPO. Such publishers would want a low reserve to achieve their goal of optimizing sold inventory. As another example, a publisher may have a data team they have invested in that sets reserve prices based on tailor-made algorithms for their user and advertiser base. As a result, this publisher may believe that their reserve generates greater revenues than what would result from RPO and prefer the outcomes with their own algorithm instead of RPO. As another possible example, a publisher might have a small data team optimizing the AdX reserve under the assumption that this is the reserve set in AdX, and that when an impression fails to sell through AdX it is because no bid above the reserve was received from AdX. This would be an incorrect assumption under RPO, because it could instead be that such a bid was received but filtered by AdX for not clearing the higher reserve set by RPO. Hence, the publisher data team might have

430 [redacted]

corrupted data, and as a result fail to optimize their reserve on AdX in the future based on this data, or mistakenly use this corrupted data to poorly set reserves for other exchanges.[431]

B. Impact of Reserve Price Optimization on Advertisers

281. RPO would lead to a payoff loss for advertisers since it leads to both a decrease in impressions won and an increase in the average price paid for impressions won. This is because the advertisers face higher reserves, which can lead to:

a. A decrease in the advertiser win rate, since RPO increases publisher reserves, and a higher reserve leads to a lower probability of the impression being cleared in a first- or second-price auction. As a result, advertisers might lose impressions that they would have won otherwise.

b. An increase in the average clearing price, during periods when AdX ran a second-price auction, since a higher reserve leads to a higher clearing price when the second highest bid is the reserve itself. Other auctions where there are two bids above the reserve either remain unaffected (if the increased reserve is still below the second highest bid) or RPO leads to a higher clearing price in these as well (if the increased reserve surpasses the second highest bid). As a result, the publishers must pay more for impressions they win compared to what they would have without RPO.

c. During periods when both RPO and Exchange Bidding were active, an increased clearing price would cause AdX to submit a higher clearing price to Exchange Bidding, which could cause AdX to win additional impressions over competing exchanges' bids. Note that this reasoning does not apply prior to Exchange Bidding. Prior to Exchange Bidding, AdX won every auction for which it found an advertiser above its reserve due to (Enhanced) Dynamic Allocation.

282. Google internal documents demonstrate this as well. A launch email states that "setting optimized prices on behalf of publishers makes queries more expensive for buyers."[432]

[431] These are all reasons why a publisher could reasonably prefer outcomes without RPO to outcomes with RPO and highlight reasons why the existence of RPO is material to publishers. These reasons also extend to any auction for which a publisher believes RPO is possible, even if RPO is ultimately inactive on that auction.
[432] GOOG-NE-06151351 at -53. November 12, 2015. Email thread, "Subject: [Monetization-pm] Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers."

1) Advertisers would change their bidding behavior had Google revealed RPO during its initial implementation

283. Google concealed vital information from advertisers by concealing RPO. In particular, if Google advertised AdX as a truthful second-price auction,[433] this would encourage advertisers to bid their true values for impressions. However, if RPO was also using past AdX bid data to set reserves for future AdX auctions, the long-term process is no longer truthful. That is, for any particular AdX auction, advertisers indeed maximize their payoff in this individual auction by reporting their true value. However, submitting a high bid equal to an advertiser's true value would cause later AdX reserves to increase, decreasing that advertiser's future payoff in later AdX auctions.[434]

284. To see why that is the case, imagine an advertiser who has just opened the first Lego shop in Plano, TX is planning to advertiser for their business. Prior to this, the demand for an impression to a user who is over the age of 25, lives in Plano, and likes Legos might be just $1, because the value to most advertisers is simply just because the user is over the age of 25 and lives in Plano. RPO would learn this over time and set a reserve of $1 for such users. This advertiser, on the other hand, has the perfect business for this user, and has a value of $10. The advertiser believes that they are participating in a regular second-price auction with reserve $1, and therefore bids $10 for all such impressions, since bidding their value is the best strategy for them in a second-price auction. Initially, they win them all and pay $1 each, because they are indeed participating in a second-price auction with reserve $1. Google's RPO algorithm notices, however, that impressions for users over the age of 25, living in Plano, and who like Legos tend to elicit bids closer to $10, which is much more than the current reserve of $1. RPO does its job and raises the reserve, eventually getting quite close to $10. This means that while the advertiser enjoyed a good profit initially, their profit for these impressions is going to converge to $0 in the long run due to RPO, since they will be paying their value for each impression.

---

[433] Google claimed that AdX ran a second-price auction up until it switched to the first-price auction format. *See* Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20240122142404/https:/blog.google/products/admanager/update-first-price-auctions-google-ad-manager/

[434] GOOG-NE-13207530 at -30. August 25, 2015. [redacted]

285. Note that each individual auction in isolation is still a second-price auction with reserve, and so an advertiser cannot improve their gain in that auction by shading their bid. However, if the advertiser were aware of RPO, they would have shaded their bid from the beginning. For example, if the advertiser instead consistently submitted a bid of $2, they would still win every impression, but they would instead have a long run per-impression profit of $8 since they value the impressions at $10 but pay $2.

C. Impact of Reserve Price Optimization on AdX

286. The impact of RPO on AdX is indeterminate and would depend on how effective RPO is. The program can lead to:

a. A decrease in the AdX win rate, since RPO increases publisher reserves and a higher reserve leads to a lower probability of the impression being cleared in a second-price auction. A lower win rate has a negative effect on AdX revenue.

b. An increase in the average AdX clearing price in the auctions they clear, since a higher reserve leads to a higher clearing price when the second highest bid is the reserve itself. Other auctions where there are two bids above the reserve either remain unaffected (if the increased reserve is still below the second highest bid) or RPO leads to a higher clearing price in these as well (if the increased reserve surpasses the second highest bid). As a result, average AdX clearing price increases, which has a positive effect on AdX revenue since AdX takes a fee based on the clearing price.

c. During periods when both RPO and Exchange Bidding were active, an increased clearing price would cause AdX to submit a higher clearing price to Exchange Bidding, which could cause AdX to win additional impressions over competing exchanges' bids. Note that this reasoning does not apply prior to Exchange Bidding. Prior to Exchange Bidding, AdX won every auction for which it found an advertiser above its reserve due to (Enhanced) Dynamic Allocation.

287. RPO is effective when the gains from b (and c, during Exchange Bidding) outweigh the losses from a, and it is reasonable to expect a sophisticated seller with ample data, such as Google, to accomplish this.

288. Google documents suggest that RPO indeed improved Google's revenue. For example, a post-launch email notes that RPO led to an increase in annual revenue [redacted] for publishers,[435] and a March 2016 brief estimates the annual increase in Google revenue due to RPO [redacted].[436]

289. RPO allows AdX to increase the reserve price beyond what the publisher sets. Dynamic Revenue Sharing[437] allows AdX to functionally lower the effective reserve price[438] below what the publisher sets.[439, 440] Both conducts together allow AdX the flexibility to adjust the publisher-set reserve in either direction.[441]

[redacted]

[437] See Section VII for a discussion of Dynamic Revenue Sharing.
[438] That is, the reserve price plus the ad exchange take rate.
[439] GOOG-NE-13205325 at -36. (comment [5] suggests further that DRSv2 could also lower the take-rate applied to an RPO-set reserve, and not just the publisher-set reserve.)
[440] In principle, DRS, if active on RPO-set reserves, can also help AdX clear impressions where RPO mistakenly sets a higher reserve price than what is needed. In such a case, DRS can help AdX clear the impression, as long as the amount that RPO is mistaken by is lower than what DRS can achieve by decreasing the AdX take rate to 0%.
[441] The magnitude with which AdX can increase the publisher's reserve using RPO is unbounded, while the magnitude with which AdX can decrease the publisher's reserve using DRS is bounded. AdX will always collect at least the publisher's price floor with a 0% take rate from the winning advertiser.

APPENDIX A. CIRRICULUM VITAE

S. Matthew Weinberg

Contact Information	231 South Stanworth Dr. Princeton, NJ 08540	*Cell:* +1 (571)-278-3990 *E-mail:* smweinberg@princeton.edu

Research Interests — Algorithmic Mechanism Design, Algorithms and Uncertainty, Mechanism Design for Blockchain Applications.

Employment

Spring 2017 - Present — **Princeton University** — **Princeton, NJ**

Associate Professor, with tenure (2023 - present).
Assistant Professor (2017 - 2023).
Instructor for COS445: Economics and Computation (Spring 2017, Spring 2018, Spring 2019, Spring 2020, Spring 2021, Spring 2022, Spring 2023, Spring 2024).
Instructor for COS521: Advanced Algorithms (Fall 2017, Fall 2019, Fall 2021, Fall 2022).
Instructor for COS597F: Open Problems in Algorithmic Game Theory (Fall 2018).
Instructor for COS597A: Advanced Topics in Computer Science: Algorithmic Mechanism Design for Cryptocurrencies and DeFi (Fall 2022).

Fall 2014 - Spring 2016 — **Princeton University** — **Princeton, NJ**

Postdoctoral Researcher.
Instructor for COS597A: Algorithmic Mechanism Design (Fall 2014).

Education

2010 - 2014 — **Massachusetts Institute of Technology** — **Cambridge, MA**

Ph.D., Computer Science

Adviser: Constantinos Daskalakis
Thesis: Algorithms for Strategic Agents
Teaching Assistant for 6.046: Design and Analysis of Algorithms (Spring 2013).
Teaching Assistant for 6.853: Topics in Algorithmic Game Theory (Fall 2011).

2006 - 2010 — **Cornell University** — **Ithaca, NY**

B.A., Mathematics

GPA: 4.038/4.3
Magna Cum Laude
Teaching Assistant for CS 4820: Introduction to Analysis of Algorithms.
Teaching Assistant for CS 2800: Discrete Structures.

Consultant for CS 2110: Object Oriented Programming and Data Structures.
Consultant for CS 100J: Introduction to Programming using Java

Grants and Awards

2022	**FOCS Test of Time Award.** For "Optimal Multi-Dimensional mechanism Design: Reducing Revenue to Welfare Maximization." Joint with Yang Cai and Constantinos Daskalakis.
2022	**President's Award for Distinguished Teaching** ("The PADT awards honor a sustained record of distinguished teaching over the course of a career at Princeton, at both undergraduate and graduate levels. Normally two senior and two junior awards are given annually." Awarded at Princeton's Class of 2022 Commencement.)
2022	**Princeton Engineering Council Lifetime Achievement Award** (awarded to faculty who have received five Engineering Council Teaching Awards).
2017, 2018, 2019, 2020, 2022	**Princeton Engineering Council Teaching Award** (awarded to ≈ 4 Engineering faculty annually, selected by students), for COS 445.
2018, 2019, 2021	**Princeton Engineering Commendation List for Distinguished Teaching**, for COS 445 (2018) and COS 521 (2019, 2021)
2020	**Sloan Foundation Fellowship.** $75,000.
2020	**Google Faculty Research Award.** $80,000. Joint with Mark Braverman at Princeton (the award of $80,000 jointly supports Mark and myself).
2020 - 2024	**NSF-1955205. Collaborative Research:AF:Medium: Moderm Combinatorial Optimization: Incentives, Uncertainty, and Smoothed Analysis.** $415,957. REU Supplement $16,000. Joint with Aviad Rubinstein at Stanford (the award of $415,957 and $16,000 REU supplement support my research at Princeton — Aviad separately received additional funding through this award).
2020 - 2024	**NSF-1942497. CAREER: Towards a Prescriptive Theory of Algorithmic Mechanism Design.** $602,818.
2019	**Phi Beta Kappa Teaching Award** (awarded annually to two Princeton professors, selected by students of Phi Beta Kappa honor society)
2019	**Howard B. Wentz Jr. Faculty Advancement Award**, from Princeton SEAS ("The award recognizes promising junior faculty members"). $50,000.
2018	**ACM Conference on Economics and Computation (EC) Best Full Paper Award** and **Best Paper with Student Lead Author Award**. For "Selling to a No-Regret Buyer." Joint with Mark Braverman, Jieming Mao, and Jon Schneider.

2017-2020	**NSF-1717899. AF: Small: Duality-based tools for simple vs. optimal mechanism design and applications to cryptocurrency**. $450,000. REU Supplement $16,000.
2014	**SIGecom Doctoral Dissertation Award**
2014	**George M. Sprowls Award** (for best MIT doctoral theses in CS)
2013	**Microsoft Research PhD Fellow**
2012	**ACM Conference on Electronic Commerce (EC) Best Paper with Student Lead Author Award**. For "Symmetries and Optimal Multi-Dimensional Mechanism Design." Joint with Constantinos Daskalakis.
2011	**National Science Foundation Graduate Research Fellow**
2010	**Akamai Presidential Fellow**
2010	**National Physical Sciences Consortium Fellow**

Graduate Advising Experience

Ariel Schvartzman (PhD 2020). Postdoc at Rutgers 2020-2022. Researcher at Google 2022 - present.

Divyarthi Mohan (PhD 2021). Postdoc at Tel Aviv University 2021 - present.

Matheus Venturyne Xavier Ferreira (PhD 2021). Postdoc at Harvard 2021 - 2023. Faculty at University of Virginia CS starting Fall 2024.

Meryem Essaidi (PhD 2022). Postdoc at UC Berkeley 2022 - present.

Clayton Thomas (PhD 2023). Postdoc at Microsoft Research 2023 - present.

Linda Cai (MSE 2020, 4th year PhD).

Qianfan Zhang (2nd year PhD).

Jingyi Liu (2nd year PhD, co-advised with Mark Braverman).

Eric Xue (2nd year PhD, co-advised with Mark Braverman).

Stephen Newman (2nd year PhD, co-advised with Mark Braverman).

Kaya Ito Alpturer (1st year PhD).

Aadityan Ganesh (1st year PhD).

Chenghan Zhou (2nd year MSE).

Hedyeh Beyhaghi (visiting PhD student from 9/2017 - 9/2019). Postdoc at Northwestern and TTI 2019 - 2021. Postdoc at CMU 2021 - present.

Sahil Singla (postdoc 2018 - 2021). Assistant Professor at Georgia Tech Computer Science 2021 - present.

Undergraduate Advising Experience

I supervised 127 advisee-semesters (90 COS, 31 MAT, 2 ORF, 1 ECE, 3 PACM) over 11 semesters from Spring 2017 - Spring 2024. Below are year-by-year details.

2023-2024

I supervised the following five (5 COS) two-semester theses: Ty Kay (COS), Nicole Klausner (COS), Ryan McDowell (COS), Jude Muriithi (COS), Dwaipayan Saha (COS), the following one (COS) two-semester project: Arya Maheshwari (COS), and the following one (ECE) one-semester independent work: Amanda Wang (ECE).

2022-2023

I supervised the following six (5 COS, 1 MAT) two-semester theses: Aditya Gollapudi (COS), Richard Huang (COS), Frederick Qiu (COS), Dwaipayan Saha (COS), Anton Stengel (MAT), Henrique Vera (COS), the following one (MAT) one-semester independent work: Dimitar Chakarov (MAT), and the following one year-long project of a recent graduate: Sumanth Maddirala (MAT 2022), and the following one year-long project of a recent graduate of CMI: Aadityan Ganesh (MAT 2022).

2021-2022

I supervised the following eight (4 COS, 4 MAT) two-semester theses: Kiril Bangachev (MAT), Atanas Dinev (MAT), Nathan Finkle (COS), Anthony Hein (COS), Liam Johansson (MAT), Rahul Saha (COS), Catherine Yu (MAT), Jerry Zhu (COS), and the following four (3 COS, 1 MAT) one-semester independent works: Emily Dale (COS), Hannah Huh (COS), Austin Li (COS), Simon Park (MAT).

2020-2021

I supervised the following six (5 COS, 1 MAT) two-semester theses: Jessica Fielding (COS), Alice Gao (COS), Tinashe Handina (COS), Victor Hua (COS), Geoffrey Mon (COS), Seyoon Ragavan (MAT), the following six (3 COS, 3 MAT) one-semester independent works: Kiril Bangachev (MAT), Atanas Dinev (MAT), Ezra Edelman (COS), Michelle Woo (COS), Catherine Yu (MAT), Noa Zarur (COS), and the following PACM certificate projects: Emily Ryu (CHM).

2019-2020

I supervised the following seven (4 COS, 3 MAT) two-semester theses: Rebecca Barber (COS), Yafah Edelman (COS), William Jiao (MAT), Jonathan Jow (COS), Kevin Lin (MAT), Tristan Pollner (MAT), Shirley Zhang (COS), the following 4 (3 COS, 1 MAT) one-semester independent works: Jacob Christensen (MAT), Kimberly Ding (COS), Frankie Lam (COS), Kawin Tiyawattanaroj (COS), and the following PACM certificate projects: Sally Hahn (MAT), Seyoon Ragavan (MAT).

2018-2019

I supervised the following ten (10 COS) two-semester theses: Maryam Bahrani (COS), Rachana Balasubramanian (COS), Natalie Collina (COS), Jonathan Jow (COS), Pelumi Odimayo (COS), Simisola Olofinboba (COS), Jose Rodriguez Quinones (COS), Mel Shu (COS), Evan Wildenhain (COS), Albert Zuo (COS), the following two (2 COS) one-semester independent works: Moyin Opeyemi (COS), Eitan Zlatin (COS), and the following summer projects: Andrei Graur (MAT), Georgy Noarov (MAT), Tristan Pollner (MAT).

2017-2018 I supervised the following six (3 COS, 2 MAT, 1 ORF) two-semester theses: Maryam Bahrani (COS), Sung Won Chang (COS), Zach Halem (ORF), Heesu Hwang (MAT), Yash Patel (MAT), Bennett Victor (COS), the following eight (5 COS, 3 MAT) one-semester independent works: Richard Adjei (COS), Hrishikesh Khandeparkar (COS), Dylan Mavrides (MAT), Eric Neyman (MAT), Terri Rossi (COS) Evan Wildenhain (COS), Daphne Yang (MAT), Jonathan Zhang (COS), and the following visiting student: Jack Wang (Harvard, COS).

2016-2017 I supervised the following two-semester thesis: Will Rose (MAT).

Publications

[QW 24] Frederick Qiu, S. Matthew Weinberg:
Settling the Communication Complexity of VCG-based Mechanisms for All Approximation Guarantees.
In the *56th Annual ACM Symposium on Theory of Computation (STOC), 2024.*

[BW 24] Maryam Bahrani, S. Matthew Weinberg:
Undetectable Selfish Mining.
In the *25th Annual ACM Conference on Economics and Computation (EC), 2024.*

[DRWX 24] Mahsa Derakhshan, Emily Ryu, S. Matthew Weinberg, Eric Xue:
Settling the Competition Complexity of Additive Buyers of Independent Items.
In the *25th Annual ACM Conference on Economics and Computation (EC), 2024.*

[BW 24] Matheus V. X. Ferreira, Aadityan Ganesh, Jack Hourigan, Hannah HuhS. Matthew Weinberg, Catherine Yu:
Computing Optimal Manipulations in Cryptographic Self-Selection Proof-of-Stake Protocols.
In the *25th Annual ACM Conference on Economics and Computation (EC), 2024.*

[CTW 24] Aadityan Ganesh, Clayton Thomas, S. Matthew Weinberg:
Revisiting the Primitives of Transaction Fee Mechanism Design.
In the *25th Annual ACM Conference on Economics and Computation (EC), 2024.*

[DW 24] Atanas Dinev, S. Matthew Weinberg:
Simple and Optimal Online Contention Resolution Schemes for k-Uniform Matroids.
In the *15th Annual Innovations of Theoretical Computer Science (ITCS), 2024.*

[CWWZ 23] Linda Cai, S. Matthew Weinberg, Evan Wildenhain, Shirley Zhang:
Selling to Multiple No-Regret Buyers.
In the *18th Conference on Web and Internet Economics (WINE), 2023.*

[CGW 23] Linda Cai, Joshua Gardner, S. Matthew Weinberg, Shirley Zhang:
Optimal Stopping with Multi-Dimensional Comparative Loss Aversion.
In the *18th Conference on Web and Internet Economics (WINE), 2023.*

[SVW 23] Raghuvansh R. Saxena, Santhoshini Velusamy, S. Matthew Weinberg:
An Improved Lower Bound for Matroid Intersection Prophet Inequalities.
In the *14th Annual Innovations of Theoretical Computer Science (ITCS), 2023.*

[DW 22] Atanas Dinev, S. Matthew Weinberg:
Tight Bounds on 3-Team Manipulations in Randomized Death Match.
In the *18th Conference on Web and Internet Economics (WINE), 2022.*

[GRW 22] Akash Gaonkar, Divya Raghunathan, S. Matthew Weinberg:
The Derby Game: An Ordering-based Colonel Blotto Game.
In the *23rd Annual ACM Conference on Economics and Computation (EC), 2022.*

[WZ 22] S. Matthew Weinberg, Zixin Zhou:
Optimal Multi-Dimensional Mechanisms are not Locally-Implementable.
In the *23rd Annual ACM Conference on Economics and Computation (EC), 2022.*

[DFRSW 22] Emily Dale, Jessica Fielding,, Hari Ramakrishnan, Sacheth Sathyanarayanan, S. Matthew Weinberg:
Approximately Strategyproof Tournament Rules with Multiple Prizes.
In the *23rd Annual ACM Conference on Economics and Computation (EC), 2022.*

[FHWY 22] Matheus V.X. Ferreira, Ye Lin Sally Hahn, S. Matthew Weinberg, Catherine Yu:
Optimal Strategic Mining against Cryptographic Self-Selection in Proof-of-Stake.
In the *23rd Annual ACM Conference on Economics and Computation (EC), 2022.*

[PSW 22] Christos-Alexandros Psomas, Ariel Schvartzman, S. Matthew Weinberg:
On Infinite Separations Between Simple and Optimal Mechanisms.
In the *35th Annual Conference and Workshop on Neural Information Processing Systems (NeurIPS), 2022.*

[CCFPW 22] Jose Correa, Andres Cristi, Andres Fielbaum, Tristan Pollner, S. Matthew Weinberg:
Optimal Item Pricing in Online Combinatorial Auctions.
In the *23rd Annual Conference on Integer Programming and Combinatorial Optimization (IPCO), 2022.*

[EFW 22] Meryem Essaidi, Matheus V. X. Ferreira, S. Matthew Weinberg:
Credible, Strategyproof, Optimal, and Bounded Expected-Round Single-Item Auctions for All Distributions.
In the *13th Annual Innovations of Theoretical Computer Science (ITCS), 2022*.

[EW 21] Meryem Essaidi, S. Matthew Weinberg:
On Symmetries in Multi-Dimensional Mechanism Design.
In the *17th Conference on Web and Internet Economics (WINE), 2021*.
Accepted for oral presentation in the *1st Annual ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO), 2021*.

[BBSW 21] Maryam Bahrani, Hedyeh Beyhaghi, Sahil Singla, S. Matthew Weinberg:
Formal Barriers to Simple Algorithms for the Matroid Secretary Problem.
In the *17th Conference on Web and Internet Economics (WINE), 2021*.

[FW 21] Matheus V. X. Ferreira, S. Matthew Weinberg:
Proof of Stake Mining Games with Perfect Randomness.
In the *22nd Annual ACM Conference on Economics and Computation (EC), 2021*.

[BSW 21] Mark Braverman, Jon Schneider, S. Matthew Weinberg:
Prior-free Dynamic Mechanism Design With Limited Liability.
In the *22nd Annual ACM Conference on Economics and Computation (EC), 2021*.

[NNW 21] Eric Neyman, Georgy Noarov S. Matthew Weinberg:
Binary Scoring Rules that Incentivize Precision.
In the *22nd Annual ACM Conference on Economics and Computation (EC), 2021*.

[RSTWZ 21] Aviad Rubinstein, Raghuvansh R. Saxena, Clayton Thomas, S. Matthew Weinberg, Junyao Zhao:
Exponential Communication Separations between Notions of Selfishness.
In the *53rd Annual ACM Symposium on Theory of Computation (STOC), 2021*.

[RJW 21] Jad Rahme, Sami Jelassi, S. Matthew Weinberg:
Auction Learning as a Two-Player Game.
In the *9th Annual International Conference on Learning Representations (ICLR), 2021*.

[RJBW 21] Jad Rahme, Sami Jelassi, Joan Bruna, S. Matthew Weinberg:
A Permutation Equivariant Neural Network Architecture for Auction Design.
In the *35th Annual AAAI Conference on Artificial Intelligence (AAAI), 2021*.

[DW 21] Kimberly Ding, S. Matthew Weinberg:
Approximately Strategyproof Tournament Rules in the Probabilistic Setting.
In the *12th Annual Innovations of Theoretical Computer Science (ITCS), 2021*.

[CKWLG 20] Michael Chang, Sidhant Kaushik, S. Matthew Weinberg, Thomas L. Griffiths, Sergey Levine
Decentralized Reinforcement Learning: Global Decision-Making via Local Economic Transactions.
In the *37th Annual International Conference on Machine Learning (ICML), 2020*.

[BIMW 20] Maryam Bahrani, Nicole Immorlica, Divyarthi Mohan, S. Matthew Weinberg:
Asynchronous Majority Dynamics in Preferential Attachment Trees.
In the *47th International Colloquium on Automata, Languages, and Programming (ICALP), 2020*.

[FW 20] Matheus V. X. Ferreira, S. Matthew Weinberg:
Credible, Truthful, and Two-Round (Optimal) Auctions via Cryptographic Commitments.
In the *21st Annual ACM Conference on Economics and Computation (EC), 2020*.

[DGSSW 20] Nikhil R. Devanur, Kira Goldner, Raghuvansh R. Saxena, Ariel Schvartzman, S. Matthew Weinberg:
Optimal Mechanism Design for Single-Minded Agents.
In the *21st Annual ACM Conference on Economics and Computation (EC), 2020*.

[CW 20] Natalie Collina, S. Matthew Weinberg:
On the (in)-approximability of Bayesian Revenue Maximization for a Combinatorial Buyer.
In the *21st Annual ACM Conference on Economics and Computation (EC), 2020*.

[AKRW20, AKRW22] Sepehr Assadi, Hrishikesh Khandeparkar, Raghuvansh R. Saxena, S. Matthew Weinberg:
Separating the Communication Complexity of Truthful and Non-Truthful Combinatorial Auctions.
In the *52nd Annual ACM Symposium on Theory of Computation (STOC), 2020*.
Accepted to Special Issue of SIAM Journal on Computing (SICOMP).

[AKW 20] Rediet Abebe, Jon Kleinberg, S. Matthew Weinberg:
Subsidy Allocations in the Presence of Income Shocks.
In the *34th Annual AAAI Conference on Artificial Intelligence (AAAI), 2020*.

[GPRW 20] Andrei Graur, Tristan Pollner, Vidhya Ramaswamy, S. Matthew Weinberg: **New Query Lower Bounds for Submodular Function Minimization**. In the *11th Annual Innovations of Theoretical Computer Science (ITCS), 2020*.

[CTW 20] Linda Cai, Clayton Thomas, S. Matthew Weinberg: **Implementation in Advised Strategies: Welfare Guarantees from Posted-Price Mechanisms when Demand Queries are NP-hard**. In the *11th Annual Innovations of Theoretical Computer Science (ITCS), 2020*.

[RWW 20] Aviad Rubinstein, Jack Z. Wang, S. Matthew Weinberg: **Optimal Single-Choice Prophet Inequalities from Samples**. In the *11th Annual Innovations of Theoretical Computer Science (ITCS), 2020*.

[SWZZ 20] Ariel Schvartzman, S. Matthew Weinberg, Eitan Zlatin, Albert Zuo: **Approximately Strategyproof Tournament Rules: On Large Manipulating Sets and Cover-Consistence**. In the *11th Annual Innovations of Theoretical Computer Science (ITCS), 2020*.

[EFNTW 19] Tomer Ezra, Michal Feldman, Eric Neyman, Inbal Talgam-Cohen, S. Matthew Weinberg: **Settling the Communication Complexity of Combinatorial Auctions with Two Subadditive Buyers**. In the *60th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2019*.

[KMSSW 19] Pravesh Kothari, Divyarthi Mohan, Ariel Schvartzman, Sahil Singla, S. Matthew Weinberg: **Approximation Schemes for a Buyer with Independent Items via Symmetries**. In the *60th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2019*.

[PSW 19] Christos-Alexandros Psomas, Ariel Schvartzman, S. Matthew Weinberg: **Smoothed Analysis of Multi-Item Auctions with Correlated Values**. In the *20th Annual ACM Conference on Economics and Computation (EC), 2019*.

[BNPW 19] Jonah Brown-Cohen, Arvind Narayanan, Christos-Alexandros Psomas, S. Matthew Weinberg: **Formal Barriers to Longest-Chain Proof-of-Stake Protocols**. In the *20th Annual ACM Conference on Economics and Computation (EC), 2019*.

[BMSW 19] Mark Braverman, Jieming Mao, Jon Schneider, S. Matthew Weinberg: **Multi-Armed Bandit Problems with Strategic Arms**. In the *32nd Annual Conference on Learning Theory (COLT), 2019*.

[BW 19] Hedyeh Beyhaghi, S. Matthew Weinberg:
Optimal (and Benchmark-Optimal) Competition Complexity for Additive Buyers over Independent Items.
In the *51st Annual ACM Symposium on Theory of Computation (STOC), 2019*.

[FWHFC 19] Matheus V. X. Ferreira, S. Matthew Weinberg, Danny Yuxing Huang, Nick Feamster, Tithi Chattopadhyay:
Selling a Single Item with Negative Externalities: to Regulate Production or Payments?.
In the *28th Annual World Wide Web Conference (WWW), 2019*.

[AW19, AW22] Nick Arnosti, S. Matthew Weinberg:
Bitcoin: A Natural Oligopoly.
In the *10th Annual Innovations of Theoretical Computer Science (ITCS), 2019*.
Accepted to Management Science.

[DNPW 19] Shaddin Dughmi, Rad Niazadeh, Christos-Alexandros Psomas, S. Matthew Weinberg:
Persuasion and Incentives through the Lens of Duality.
In the *15th Annual Conference on Web and Internet Economics (WINE), 2019*.

[GW18, GW21] Yannai Gonczarowski, S. Matthew Weinberg:
The Sample Complexity of up-to-ε Multi-Dimensional Revenue Maximization.
In the *59th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2018*.
Accepted to Journal of the ACM (JACM).

[BMSW 18] Mark Braverman and Jieming Mao and Jon Schneider and S. Matthew Weinberg:
Selling to a No-Regret Buyer.
In the *19th ACM Conference on Economics and Computation (EC), 2018*.
Best Full Paper Award.
Best Paper with Student Lead Author Award.

[KGCWF 18] Harry Kalodner, Steven Goldfeder, Xiaoqi Chen, S. Matthew Weinberg, Edward W. Felten:
Arbitrum: Scalable Smart Contracts.
In the *27th Annual USENIX Security Symposium (USENIX), 2018*.

[RSW 18] Aviad Rubinstein, Tselil Schramm, S. Matthew Weinberg:
Computing exact minimum cuts without knowing the graph.
In the *9th Annual Innovations of Theoretical Computer Science (ITCS), 2018*.

[BMW 18] Mark Braverman, Jieming Mao, S. Matthew Weinberg:
On Simultaneous Two-Player Combinatorial Auctions.

In the *29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2018.*

[SSW 18] Raghuvansh R. Saxena, Ariel Schvartzman, S. Matthew Weinberg:
The Menu Complexity of "one-and-a-half-dimensional" mechanism design.
In the *29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2018.*

[DW 17] Nikhil R. Devanur, S. Matthew Weinberg:
The Optimal Mechanism for Selling to a Budget Constrained Buyer: The General Case.
In the *18th Annual ACM Conference on Economics and Computation (EC), 2017.*

[EFFTW17b, EFFTW21] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, S. Matthew Weinberg:
A Simple and Approximately Optimal Mechanism for a Buyer with Complements.
In the *18th Annual ACM Conference on Economics and Computation (EC), 2017.*
Accepted to Operations Research.

[EFFTW 17a] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, S. Matthew Weinberg:
The Competition Complexity of Auctions: A Bulow-Klemperer Result for Multi-Dimensional Bidders.
In the *18th Annual ACM Conference on Economics and Computation (EC), 2017.*

[HWZJC 17] Zhe Huang, S. Matthew Weinberg, Liang Zheng, Carlee Joe-Wong, Mung Chiang:
Discovering Valuations and Enforcing Truthfulness in a Deadline-Aware Scheduler.
In the *37th Annual IEEE Conference on Computer Communications (INFOCOM), 2017.*

[SSW 17] Jon Schneider, Ariel Schvartzman, S. Matthew Weinberg:
Condorcet-Consistent and Approximately Strategyproof Tournament Rules.
In the *8th Innovations of Theoretical Computer Science Conference (ITCS), 2017.*

[CKWN 16] Miles Carlsten, Harry Kalodner, S. Matthew Weinberg, Arvind Narayanan:
On the Instability of Bitcoin without the Block Reward.
In the *23rd Annual ACM Conference on Computer and Communications Security (CCS), 2016.*

[CDW16, CDW21] Yang Cai, Nikhil Devanur, S. Matthew Weinberg:
A Duality Based Unified Approach to Bayesian Mechanism Design.
In the *48th Annual ACM Symposium on Theory of Computation (STOC),*

2016.
Accepted to Special Issue of SIAM Journal on Computing (SICOMP).

[BMW 16b] Mark Braverman, Jieming Mao, S. Matthew Weinberg:
Parallel Algorithms for Select and Partition with Noisy Comparisons.
In the *48th Annual ACM Symposium on Theory of Computation (STOC), 2016*.

[BMW 16a] Mark Braverman, Jieming Mao, S. Matthew Weinberg:
Interpolating Between Truthful and Non-truthful Mechanisms for Combinatorial Auctions.
In the *27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2016*.

[RW 15, RW 18] Aviad Rubinstein, S. Matthew Weinberg:
Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity.
In the *16th Annual ACM Conference on Economics and Computation (EC), 2015*.
Accepted to Special Issue of Transactions on Economics and Computation (TEAC).

[DDW15,DDW18] Constantinos Daskalakis, Nikhil Devanur, S. Matthew Weinberg:
Revenue Maximization and Ex-Post Budget Constraints.
In the *16th Annual ACM Conference on Economics and Computation (EC), 2015*.
Accepted to Special Issue of Transactions on Economics and Computation (TEAC).

[DMSW 15] Nikhil Devanur, Jamie Morgenstern, Vasilis Syrgkanis, S. Matthew Weinberg:
Simple Mechanisms with Simple Strategies.
In the *16th Annual ACM Conference on Economics and Computation (EC), 2015*.

[DW 15] Constantinos Daskalakis, S. Matthew Weinberg:
Bayesian Truthful Mechanisms for Job Scheduling from Bi-criterion Approximation Algorithms.
In the *26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2015*.

[DKTWW 15] Constantinos Daskalakis, Nicolaas Kaashoek, Christos Tzamos, S. Matthew Weinberg, William Wu:
Game Theory Based Peer Grading for MOOCs.
In the *Second ACM Conference on Learning at Scale (L@S), 2015*.
Work in Progress paper.

[BILW14, BILW20] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, S. Matthew Weinberg:
A Simple and Approximately Optimal Mechanism for an Additive Buyer.

In the *55th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2014.*
Accepted to Journal of the ACM (JACM).

[FILW 14] Michal Feldman, Nicole Immorlica, Brendan Lucier, S. Matthew Weinberg:
Reaching Consensus via non-Bayesian Asynchronous Learning in Social Networks.
In the *17th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX), 2014.*

[AKW14, AKW19] Pablo D. Azar, Robert Kleinberg, S. Matthew Weinberg:
Prophet Inequalities with Limited Information.
In the *25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2014.*
Accepted to Special Issue of Games and Economic Behavior (GEB).

[CDW 13b] Yang Cai, Constantinos Daskalakis and S. Matthew Weinberg:
Understanding Incentives: Mechanism Design Becomes Algorithm Design.
In the *54th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2013.*

[CDW 13a] Yang Cai, Constantinos Daskalakis and S. Matthew Weinberg:
Reducing Revenue to Welfare Maximization: Approximation Algorithms and other Generalizations.
In the *24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2013.*

[ADMW 13] Pablo Azar, Constantinos Daskalakis, Silvio Micali and S. Matthew Weinberg:
Optimal and Efficient Parametric Auctions.
In the *24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2013.*

[CDW 12b] Yang Cai, Constantinos Daskalakis and S. Matthew Weinberg:
Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization.
In the *53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2012.*
FOCS 2022 Test of Time Award.

[DW 12] Constantinos Daskalakis and S. Matthew Weinberg:
Symmetries and Optimal Multi-Dimensional Mechanism Design.
In the *13th ACM Conference on Electronic Commerce (EC), 2012.*
Best Paper with Student Lead Author Award.

[CDW 12a] Yang Cai, Constantinos Daskalakis and S. Matthew Weinberg:
An Algorithmic Characterization of Multi-Dimensional Mechanisms.
In the *44th ACM Symposium on Theory of Computing (STOC), 2012.*

[KW 12, KW 19] Robert Kleinberg and S. Matthew Weinberg:
Matroid Prophet Inequalities.
In the *44th ACM Symposium on Theory of Computing (STOC), 2012.*
Accepted to Special Issue of Games and Economic Behavior (GEB).

[BCKW10,BCKW15] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg:
Pricing Randomized Allocations.
In the *21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2010.*
Accepted to Special Issue of Journal of Economic Theory (JET).

Organized Workshops and Tutorials

SIGecom Winter Meeting on Web3/Blockchain/Cryptocurrencies (with Scott Kominers), 2023.

Economics of Distributed Systems (with Jacob Leshno), Conference on Economics and Computation, 2022. Conference on Web and Internet Economics, 2022.

Algorithmic Game Theory Mentoring Workshop (with Michal Feldman and Raf Frongillo), Conference on Economics and Computation, 2020.

Early Career TCS Mentoring Workshop (with Yael Kalai and Aviad Rubinstein), Symposium on the Foundations of Computer Science, 2019.

Algorithmic Game Theory Mentoring Workshop (with Ariel Procaccia and Daniela Saban), Conference on Economics and Computation, 2019.

Tutorial on Incentives in Cryptocurrencies (with Jacob Leshno, Arvind Narayanan, Georgios Piliouras and Christos-Alexandros Psomas), Conference on Economics and Computation, 2018.

Algorithmic Game Theory Mentoring Workshop (with Nicole Immorlica and Ruta Mehta), Conference on Economics and Computation, 2018.

Workshop on Connections between Theory of Computation and Mechanism design (with Yang Cai and Shuchi Chawla), Symposium on Theory of Computing, 2017.

Tutorial on Prophet Inequalities and Secretary Problems, Simons Institute Bootcamp for Algorithms and Uncertainty, Berkeley, 2016.

Tutorial on Bayesian Mechanism Design (with Yang Cai and Constantinos Daskalakis), Conference on Economics and Computation, 2014.

Professional Service

ACM Transactions on Economics and Computation (TEAC): Associate Editor (2022 - present).

ACM SIGecom Exchanges: Co-Editor (2018-2022).

ACM Symposium on the Theory of Computing (STOC): Program Committee (2022, 2024).

ACM Conference on Economics and Computation (EC): Area Chair (2021, 2023), Senior Program Committee (2017, 2019, 2020, 2024), Treasurer (2016), Program Committee (2015, 2016, 2018, 2022). Distinguished Senior PC member (2019). Distinguished PC member (2018).

ACM-SIAM Symposium on Discrete Algorithms (SODA): Program Committee (2017, 2022)

Advances in Financial Technologies (AFT): Program Chair (2023), General Chair (2023), Program Committee (2019, 2020, 2021)

SIAM Journal on Computing (SICOMP): Associate Editor for STOC 2024 Special Issue (2024).

International Colloquium on Automata, Languages and Programming (ICALP): Program Committee (2017)

Innovations in Theoretical Computer Science (ITCS): Program Committee (2019, 2021, 2024)

ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO): Area Chair (2021)

International World Wide Web Conference (WWW): Program Committee (2017), Outstanding Reviewer Award (2017)

Conference on Web and Internet Economics (WINE): Senior Program Committee (2022), Program Committee (2016, 2017, 2018, 2019, 2020)

Workshop on the Economics of Networks, Systems and Computation (NetEcon): Program Committee (2016, 2018)

Conference on Auctions, Market Mechanisms and Their Applications (AMMA): Program Committee (2015)

International Conference on Artificial Intelligence (IJCAI): Program Committee (2013)

IEEE International Workshop on Network Science for Communication Networks: Program Committee (2017)

Workshop on Mechanism Design for Social Good (MD4SG): Program Committee (2017, 2018, 2019, 2020)

Computing Research Association (CRA): Undergraduate Research Awards Committee Member (2022, 2023).

Departmental and University Service

SEAS DeCenter. Associate Director (07/2023 - present), Steering Committee (06/2022 - present), Faculty Search Committee (2022-2023, 2023-2024).

Council of the Princeton University Community. Member (06/2023 - present), Executive Committee (06/2023 - present), Faculty Advisory Committee on Policy (06/2023 - present).

COS Committee on Diversity, Climate, and Inclusion, Member (09/2018 - present), Lead for issues related to grad admissions (01/2021 - 08/2022), Chair (09/2022 - present).

Center for Information and Technology Policy. Faculty Search Committee (2023-2024).

Council on Teaching and Learning, Member (09/2020 - 06/2023).

Faculty advisor for JuST (2020 - 2022).

Theory group lead for 2019-2020 and 2021-2022 (co-led with Huacheng Yu) graduate admissions.

Co-coordinated (with Adam Finkelstein) COS TA assignment for fall 2021 and spring 2022.

Mathey BSE Advisor (09/2017 - 06/2018). COS BSE Academic Advisor (90/2018 - 6/2021). Whitman AB Advisor (09/2021 - present).

Undergraduate advising: 101 advisee-semesters (68 COS, 28 MAT, 2 ORF, 3 PACM) from 01/2017 - 06/2022.

President's Award for Distinguished Teaching. Selection Committee (2023).

Visiting Positions and Internships

Summer 2022 — **a16z Crypto** — **New York, NY**
Summer Faculty Fellow.

Fall 2016 — **Simons Institute for the Theory of Computing** — **Berkeley, CA**
Research Fellow, Semester on Algorithms and Uncertainty.

Fall 2015 — **Simons Institute for the Theory of Computing** — **Berkeley, CA**
Research Fellow, Semester on Economics and Computation.

Summer 2013 — **Microsoft Research - New England** — **Cambridge, MA**
Research Intern with Nicole Immorlica and Brendan Lucier.

Summer 2011 — **Department of Defense** — **Fort Meade, MD**
NPSC Intern. References available upon request.

Summer 2010 — **Institute for Defense Analyses** — **Princeton, NJ**
SCAMP Participant. References available upon request.

Summer 2009	**Department of Defense**	**Fort Meade, MD**
	Director's Summer Program. References available upon request.	
Summer 2008	**University of Maryland**	**College Park, MD**
	REU in network security with Professor Michel Cukier.	

APPENDIX B. MATERIALS RELIED UPON & MATERIALS CONSIDERED

MATERIALS RELIED UPON

Bates Stamped Documents

1. GOOG-AT-MDL-000875073. August 2019. "The Unified First Price Auction."
2. GOOG-AT-MDL-000987708. April 9, 2021. "PM Perspective on 1P RPO."
3. GOOG-AT-MDL-001004706. "Ad Manager Ecosystem 101."
4. GOOG-AT-MDL-001412616. "Project Bernanke and margins story."
5. GOOG-AT-MDL-001811992. June 2017. "Exchange Bidding / Platform StratOps Meeting."
6. GOOG-AT-MDL-003849201. Email thread, "Subject: Re: Partner revshare scaling factor alerts."
7. GOOG-AT-MDL-004016180. February 20, 2020. "Auction Theory Primer."
8. GOOG-AT-MDL-004017152.
9. GOOG-AT-MDL-006218257. December 16, 2022. "Case AT.40670 - Google - Adtech and Data-related practices."
10. GOOG-AT-MDL-008842383. August 5, 2023. "Declaration of Nirmal Jayaram."
11. GOOG-AT-MDL-008842393. August 4, 2023. "Declaration of Nitish Korula."
12. GOOG-AT-MDL-008881638. October 30, 2014. "Rethinking Bernanke: Grid search to line search."
13. GOOG-AT-MDL-008991406.
14. GOOG-AT-MDL-019244499. "Truthful DRS Auction Walkthrough."
15. GOOG-DOJ-15769995. May 2017. "Protecting Publishers from Objectionable Ads - Proposal."
16. GOOG-DOJ-27769247. September 2, 2016. "Header Bidding and FAN."
17. GOOG-DOJ-28385887. August 17, 2015. "Beyond Bernanke."
18. GOOG-DOJ-28386151. December 10, 2013. "Project Bernanke - Quantitative Easing on the AdExchange."
19. GOOG-DOJ-29803801. March 2016. "RPO brief."
20. GOOG-DOJ-AT-01809483. March 2017. "Exchange Bidding in Dynamic Allocation (fka Project Jedi)."
21. GOOG-DOJ-AT-01815211. October 2019. "Open Bidding (fka Exchange Bidding)Training."
22. GOOG-DOJ-AT-02224828. March 2019. "The Alchemist."
23. GOOG-DOJ-AT-02471194. July 26, 2015. "Global Bernanke."
24. GOOG-DOJ-AT-02513569. "gTrade Team Background."
25. GOOG-DOJ-AT-02639830. April 2016. "Exchange Bidding (aka Jedi) LPS AMERICAS TRAINING."
26. GOOG-NE-03597611. December 15, 2011. "Mysteries of Dynamic Allocation."
27. GOOG-NE-03640022. "AdX Managed Reserves."
28. GOOG-NE-03872763. "Discussion on improving AdX & AdSense backfill."
29. GOOG-NE-03995243. July 25, 2018. "PRD: Unified 1P auction and Pricing rules."
30. GOOG-NE-04934281. July 30, 2018. "Dynamic Revenue Share."
31. GOOG-NE-05279363. "Bidding in adversarial auctions."
32. GOOG-NE-06151351. November 12, 2015. Email thread, "Subject: [Monetization-pm] Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers."
33. GOOG-NE-06839089. "Project Bernanke."
34. GOOG-NE-06842715. May 10, 2016. "AdX Auction Optimizations."

35. GOOG-NE-06864639. May 9, 2014. "Dynamic Sell-side Revshare on AdX."
36. GOOG-NE-08112779. "PBS Basics Training (3) AdX Basics."
37. GOOG-NE-09485306. December 18, 2017. "OLD – New Ad Manager Indirect Notes."
38. GOOG-NE-10780865. May 5, 2020. "Clearing Up Misconceptions About Google's Ad Tech Business."
39. GOOG-NE-11753797. February 11, 2019. "DVAA Quality, Formats, O&O - Q1 2019 All Hands."
40. GOOG-NE-12737317. Email thread, "Subject: [Follow up] DRX Suite Commercialization (11/13/14)."
41. GOOG-NE-13200831. "The case for encouraging buyers to declare two bids."
42. GOOG-NE-13203009. "DRX Global Optimization of DRS, RPO, and EDA."
43. GOOG-NE-13204729. August 17, 2015. "AdX Dynamic Price"
44. GOOG-NE-13205325.
45. GOOG-NE-13207241. "AdX Dynamic Revshare v2: Launch Doc."
46. GOOG-NE-13207530. [redacted]
47. GOOG-NE-13214748. "Modeling Design Doc for Truthful DRS."
48. GOOG-NE-13222752. August 31, 2018. "Make Combined Auction Side Effect Free and Rewindable."
49. GOOG-NE-13226622. "Truthful DRS Design Doc."
50. GOOG-NE-13234466. "Overall Pub Yield With DRS(v2)."
51. GOOG-NE-13468541. "Bernanke experiment analysis."
52. GOOG-NE-13494966. May 2019. "Managing Yield."
53. GOOG-TEX-00000744. April 26, 2017. "Exchange Bidding (Jedi) Open Beta Sates Readiness Review."
54. GOOG-TEX-00105361. April 28, 2017. "FAN Bidding in to DPI and AdMob."
55. GOOG-TEX-00594205. "The Unified First Price Auction Best Practices."
56. GOOG-TEX-00777528. Email thread, "Subject: Re: [Monetization-pm] Re: [drx-pm] LAUNCHED! AdX Dynamic Revenue Share (DRS)."
57. GOOG-TEX-00831090. April 17, 2017. "DRX 2.0 Quality."
58. GOOG-TEX-00841386. "Adx First Price Auction."
59. GOOG-TEX-00843142. September 3, 2019. "First-price bidding Update."
60. GOOG-TEX-00858434. January 29, 2020. "Dynamic Revenue Share."

Public Materials

1. AdButler. "Get the Industry-Leading Ad Server." Accessed on May 31, 2024. https://web.archive.org/web/20231129061555/https:/www.adbutler.com/pricing.html
2. AdExchanger. "Google's Exchange Bidding Is Now 'Open Bidding'; Market Researchers Slip" (August 27, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20220523024855/https:/www.adexchanger.com/ad-exchange-news/tuesday-27082019/
3. AdGlare. "Plans & Pricing." Accessed on May 31, 2024. https://web.archive.org/web/20231203094651/https:/www.adglare.com/pricing
4. Alaei, Hartline, Niazadeh, Pountourakis, and Yuan. "Optimal auctions vs. anonymous pricing." *Games and Economic Behavior* vol. 118. 2019. pg. 494-510.
5. Aleksandrs Slivkins. "Introduction to Multi-Armed Bandits" (November 2019). https://arxiv.org/pdf/1904.07272

6. Anthony Vargas. “AdExplainer: Client-Side vs. Server-Side Header Bidding: What’s The Difference?” (December 1, 2023). Accessed on May 31, 2024. https://web.archive.org/web/20240314163210/https:/www.adexchanger.com/adexplainer/adexplainer-client-side-vs-server-side-header-bidding-whats-the-difference/
7. Asmaâ Bentahar. “Bid Adjustments Simplified: Run Fair Auctions with no Hassle” (May 2, 2021). Accessed on May 31, 2024. https://web.archive.org/web/20231202021004/https:/www.pubstack.io/topics/bid-adjustments-simplified
8. Bernard Lebrun. “Existence of an Equilibrium in First Price Auctions.” *Economic Theory* vol. 7, no. 3. 1996. pg. 421–443.
9. Bernard Lebrun. “First Price Auctions in the Asymmetric N Bidder Case.” *International Economic Review* vol. 40, no. 1. 1999. pg. 125–142.
10. Bernard Lebrun. “Uniqueness of the equilibrium in first-price auctions.” *Games and Economic Behavior* vol. 55, no. 1. 2006. pg. 131-151.
11. Google. “About line items.” Accessed on May 31, 2024. https://web.archive.org/web/20240216155011/https:/support.google.com/admanager/answer/9405477?hl=en
12. Google. “Ad Exchange line items.” Accessed on May 31, 2024. https://web.archive.org/web/20221012075051/https:/support.google.com/admanager/answer/188523?hl=en
13. Google. “Ad Manager payment timelines.” Accessed on May 31, 2024. https://web.archive.org/web/20221209003032/https:/support.google.com/admanager/answer/2671030?hl=en
14. Google. “Determine a bid strategy based on your goals.” Accessed on June 6, 2024. https://web.archive.org/web/20240602100502/https://support.google.com/google-ads/answer/2472725
15. Google. “DoubleClick Advertising Exchange.” Accessed on May 31, 2024. https://web.archive.org/web/20071001100309/http:/www.doubleclick.com/products/advertisingexchange/index.aspx
16. Google. “Line item types and priorities.” Accessed on May 31, 2024. https://web.archive.org/web/20240216154938/https:/support.google.com/admanager/answer/177279?hl=en
17. Google. “Maximizing advertising revenues for online publishers.” Accessed on May 31, 2024. https://web.archive.org/web/20160911040651/https:/static.googleusercontent.com/media/www.google.com/en/googleblogs/pdfs/revenue_maximization_090210.pdf
18. Google. “Price Priority line items.” Accessed on May 31, 2024. https://web.archive.org/web/20240216154933/https:/support.google.com/admanager/answer/79306?hl=en
19. Google. “Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers.” Accessed on May 31, 2024. https://web.archive.org/web/20120130063019/http:/static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en/us/doubleclick/pdfs/DC_Ad_Exchange_WP_100713.pdf
20. Google. “Sponsorship line items.” Accessed on May 31, 2024. https://web.archive.org/web/20221209041446/https:/support.google.com/admanager/answer/177426?hl=en
21. Google. “Understand Direct and Programmatic Ad Revenue.” Accessed on May 31, 2024. https://web.archive.org/web/20231226200704/https:/newsinitiative.withgoogle.com/resources/trainings/grow-digital-ad-revenue/understand-direct-and-programmatic-ad-revenue/

22. Google. “Unified pricing rules.” Accessed on May 31, 2024. https://web.archive.org/web/20230208153751/https:/support.google.com/admanager/answer/9298008?hl=en
23. Google. “Value CPM.” Accessed on May 31, 2024. https://web.archive.org/web/20221202071803/https:/support.google.com/admanager/answer/177222?hl=en
24. Google’s First Am. Resps. and Objs. to Plaintiff’s Third Set of Interrogs. (May 24, 2024) at 11.
25. Hartline and Roughgarden. “Simple versus Optimal Mechanisms.” *Proceedings of the 10th ACM Conference on Electronic Commerce*. 2009. pg. 225-234.
26. Interactive Advertising Bureau. “OpenRTB.” Accessed on June 4, 2024. https://web.archive.org/web/20240326073202/https://www.iab.com/guidelines/openrtb/
27. Jacques Crémer and Richard P. McLean. “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions.” *Econometrica* vol. 56, no. 6. 1988. pg. 1247–1257.
28. Jason Bigler. “An update on first price auctions for Google Ad Manager” (May 10, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20240122142404/https:/blog.google/products/admanager/update-first-price-auctions-google-ad-manager/
29. Jin, Lu, Qi, Tang, and Xiao. “Tight Approximation Ratio of Anonymous Pricing.” *Proceedings of the 51st Annual ACM SIGACT Symposium on the Theory of Computing*. 2019. pg. 674-685.
30. Jonathan Bellack. “Improving yield, speed and control with DoubleClick for Publishers First Look and exchange bidding” (April 13, 2016). Accessed on May 31, 2024. https://web.archive.org/web/20240206070512/https:/blog.google/products/admanager/improving-yield-speed-and-control-with-dfp-first-look-and-exchange-bidding/
31. Jonathan Bellack. “Introducing Google Ad Manager” (June 27, 2018). Accessed on May 31, 2024. https://web.archive.org/web/20240112234145/https:/blog.google/products/admanager/introducing-google-ad-manager/
32. Jonathan Bellack. “Smarter optimizations to support a healthier programmatic market” (May 12, 2016). Accessed on May 31, 2024. https://web.archive.org/web/20200929015943/https:/blog.google/products/admanager/smarter-optimizations-to-suppor/
33. Jean-Jacques Laffont and Jacques Robert. “Optimal auction with financially constrained buyers.” *Economics Letters* vol. 52, no. 2. 1996. pg. 181-186.
34. Martin L. Weitzman. “Optimal Search for the Best Alternative.” *Econometrica* vol. 47, no. 3. 1979. pg. 641–654.
35. Mohammad Akbarpour and Shengwu Li. “Credible Auctions: A Trilemma.” *Econometrica* vol. 88, no. 2. 2020. pg. 425-467.
36. News Corp Australia. “Submission To the Australian Competition and Consumer Commission” (May 2020). Accessed on May 31, 2024. https://web.archive.org/web/20221012074940/http:/www.accc.gov.au/system/files/News%20Corp%20Australia%20(15%20May%202020).pdf
37. NT Technology. “Why Is Targeting in Programmatic Ads Better Than Usual?” Accessed on June 5, 2024. https://web.archive.org/web/20240228203217/https://nt.technology/en/faq-en/why-is-targeting-in-programmatic-ads-better-than-usual/
38. Oscar Bole. “How to make sure your floor setup is flawless?” (June 27, 2022). Accessed on May 31, 2024. https://web.archive.org/web/20231202022909/https:/www.pubstack.io/topics/implementing-floor-prices#1

39. Paul R. Milgrom and Robert J. Weber. "A Theory of Auctions and Competitive Bidding." *Econometrica* vol. 50, no. 5. 1982. pg. 1089–1122.
40. Prebid. "Boost Programmatic Advertising Revenue." Accessed on June 3, 2024. https://prebid.org/
41. Robert B. Wilson. "Competitive Bidding with Disparate Information." *Management Science* vol. 15, no. 7. 1969. pg. 446–448.
42. Roger B. Myerson. "Incentive Compatibility and the Bargaining Problem." *Econometrica* vol. 47, no. 1. 1979. pg. 61–73.
43. Roger B. Myerson. "Optimal Auction Design." *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73.
44. Sarah Sluis. "Big Changes Coming To Auctions, As Exchanges Roll The Dice On First-Price" (September 5, 2017). Accessed on May 31, 2024. https://web.archive.org/web/20220712083559/https:/www.adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price/
45. Sarah Sluis. "Google Switches To First-Price Auction" (March 6, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20220910040643/https:/www.adexchanger.com/online-advertising/google-switches-to-first-price-auction/
46. Sarah Sluis. "Publishers Lash Out Against Google Over 'Unified Pricing' Changes" (April 18, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20221102060335/https:/www.adexchanger.com/online-advertising/publishers-lash-out-against-google-over-unified-pricing-changes/
47. Sarah Sluis. "The Rise Of 'Header Bidding' And The End Of The Publisher Waterfall" (June 18, 2015). Accessed on May 31, 2024. https://web.archive.org/web/20240216154913/https:/www.adexchanger.com/publishers/the-rise-of-header-bidding-and-the-end-of-the-publisher-waterfall/
48. Shengwu Li. "Obviously Strategy-Proof Mechanisms." *American Economic Review* vol. 107, no. 11. 2017. pg. 3257-3287.
49. Sleator and Tarjan. "Amortized efficiency of list update and paging rules." *Communications of the ACM* vol. 28, no. 2. 1985. pg. 202-208.
50. Krengel and Sucheston. "On semiamarts, amarts, and processes with finite value." *Probability on Banach spaces*. 1978. pg. 197-266.
51. Vijay Krishna. "Auction Theory" (2009). Academic Press. (Chapter 1.)
52. William Vickrey. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *The Journal of Finance* vol. 16, no. 1. 1961. pg. 8-37.

MATERIALS CONSIDERED

Discovery Responses

All available discovery responses produced within the matter of *The State of Texas, et al. v. Google*, Case Number: 4:20-cv-00957-SDJ, including:

1. The Parties' amended initial disclosures;

2. The Parties' discovery responses and objections to Interrogatories, Requests for Admission, and Requests for Production; and
3. Google's written responses to Plaintiffs' Rule 30(b)(6) Notice.

<u>Deposition Transcripts & Exhibits</u>

All available deposition transcripts and exhibits within the matter of *The State of Texas, et al. v. Google*, Case Number: 4:20-cv-00957-SDJ, including:

1. Deposition and Exhibits of [redacted]
2. Deposition and Exhibits of [redacted]
3. Deposition and Exhibits of [redacted]
4. Deposition and Exhibits of [redacted]
5. Deposition and Exhibits of [redacted]
6. Deposition and Exhibits of [redacted]
7. Deposition and Exhibits of [redacted]
8. Deposition and Exhibits of [redacted]
9. Deposition and Exhibits of [redacted]
10. Deposition and Exhibits of [redacted]
11. Deposition and Exhibits of [redacted]
12. Deposition and Exhibits of [redacted]
13. Deposition and Exhibits of [redacted]
14. Deposition and Exhibits of [redacted]
15. Deposition and Exhibits of [redacted]
16. Deposition and Exhibits of [redacted]
17. Deposition and Exhibits of [redacted]
18. Deposition and Exhibits of [redacted]
19. Deposition and Exhibits of [redacted]
20. Deposition and Exhibits of [redacted]
21. Deposition and Exhibits of [redacted]
22. Deposition and Exhibits of [redacted]
23. Deposition and Exhibits of [redacted]
24. Deposition and Exhibits of [redacted]
25. Deposition and Exhibits of [redacted]
26. Deposition and Exhibits of [redacted]
27. Deposition and Exhibits of [redacted]
28. Deposition and Exhibits of [redacted]
29. Deposition and Exhibits of [redacted]
30. Deposition and Exhibits of [redacted]
31. Deposition and Exhibits of [redacted]

[redacted]

32. Deposition and Exhibits of [redacted]
33. Deposition and Exhibits of [redacted]
34. Deposition and Exhibits of [redacted]
35. Deposition and Exhibits of [redacted]
36. Deposition and Exhibits of [redacted]
37. Deposition and Exhibits of [redacted]
38. Deposition and Exhibits of [redacted]
39. Deposition and Exhibits of [redacted]
40. Deposition and Exhibits of [redacted]
41. Deposition and Exhibits of [redacted]
42. Deposition and Exhibits of [redacted]
43. Deposition and Exhibits of [redacted]
44. Deposition and Exhibits of [redacted]
45. Deposition and Exhibits of [redacted] 2024
46. Deposition and Exhibits of [redacted]
47. Deposition and Exhibits of [redacted]
48. Deposition and Exhibits of [redacted]
49. Deposition and Exhibits of [redacted]
50. Deposition and Exhibits of [redacted]
51. Deposition and Exhibits of [redacted]
52. Deposition and Exhibits of [redacted]
53. Deposition and Exhibits of [redacted]
54. Deposition and Exhibits of [redacted]
55. Deposition and Exhibits of [redacted]
56. Deposition and Exhibits of [redacted]
57. Deposition and Exhibits of [redacted]
58. Deposition and Exhibits of [redacted]
59. Deposition and Exhibits of [redacted]
60. Deposition and Exhibits of [redacted]
61. Deposition and Exhibits of [redacted]
62. Deposition and Exhibits of [redacted]
63. Deposition and Exhibits of [redacted]
64. Deposition and Exhibits of [redacted]
65. Deposition and Exhibits of [redacted]
66. Deposition and Exhibits of [redacted]
67. Deposition and Exhibits of [redacted]
68. Deposition and Exhibits of [redacted]
69. Deposition and Exhibits of [redacted]
70. Deposition and Exhibits of [redacted]
71. Deposition and Exhibits of [redacted]

[redacted]

72. Deposition and Exhibits of [redacted]
73. Deposition and Exhibits of [redacted]
74. Deposition and Exhibits of [redacted] Kratkiewicz), May 1, 2024
75. Deposition and Exhibits of [redacted] Puckett), May 1, 2024
76. Deposition and Exhibits of [redacted] Kuchta), May 3, 2024
77. Deposition and Exhibits of [redacted]
78. Deposition and Exhibits of [redacted]
79. Deposition and Exhibits of [redacted]
80. Deposition and Exhibits of [redacted]
81. Deposition and Exhibits of [redacted]
82. Deposition and Exhibits of [redacted] April 29, 2024
83. Deposition and Exhibits of [redacted]
84. Deposition and Exhibits of [redacted]

All available deposition transcripts and exhibits within the matter of *USA v. Google*, Case Number: 1:23-cv-00108-LMB-JFA, including:

[redacted]

85. Deposition and Exhibits of [redacted]
86. Deposition and Exhibits of [redacted]
87. Deposition and Exhibits of [redacted]
88. Deposition and Exhibits of [redacted]
89. Deposition and Exhibits of [redacted]
90. Deposition and Exhibits of [redacted]
91. Deposition and Exhibits of [redacted]
92. Deposition and Exhibits of [redacted]
93. Deposition and Exhibits of [redacted]
94. Deposition and Exhibits of [redacted]
95. Deposition and Exhibits of [redacted]
96. Deposition and Exhibits of [redacted]
97. Deposition and Exhibits of [redacted]
98. Deposition and Exhibits of [redacted]
99. Deposition and Exhibits of [redacted]
100. Deposition and Exhibits of [redacted]
101. Deposition and Exhibits of [redacted]

[redacted]

102. Deposition and Exhibits of [redacted]
103. Deposition and Exhibits of [redacted]
104. Deposition and Exhibits of [redacted]
105. Deposition and Exhibits of [redacted]
106. Deposition and Exhibits of [redacted]
107. Deposition and Exhibits of [redacted]
108. Deposition and Exhibits of [redacted]
109. Deposition and Exhibits of [redacted]

110. Deposition and Exhibits o[redacted]
111. Deposition and Exhibits o[redacted]
112. Deposition and Exhibits o[redacted]
113. Deposition and Exhibits o[redacted]
114. Deposition and Exhibits o[redacted]
115. Deposition and Exhibits o[redacted]
116. Deposition and Exhibits o[redacted]
117. Deposition and Exhibits o[redacted]
118. Deposition and Exhibits o[redacted]
119. Deposition and Exhibits o[redacted]
120. Deposition and Exhibits o[redacted]
121. Deposition and Exhibits o[redacted]
122. Deposition and Exhibits o[redacted]
123. Deposition and Exhibits o[redacted]
124. Deposition and Exhibits o[redacted]
125. Deposition and Exhibits o[redacted]
126. Deposition and Exhibits o[redacted]
127. Deposition and Exhibits o[redacted]
128. Deposition and Exhibits o[redacted]
129. Deposition and Exhibits o[redacted]
130. Deposition and Exhibits o[redacted]
131. Deposition and Exhibits o[redacted]
132. Deposition and Exhibits o[redacted]
133. Deposition and Exhibits o[redacted]
134. Deposition and Exhibits o[redacted]
135. Deposition and Exhibits o[redacted]
136. Deposition and Exhibits o[redacted]
137. Deposition and Exhibits o[redacted]
138. Deposition and Exhibits o[redacted]

Expert Reports & Declarations

All available expert reports (with redactions) within the matter of *USA v. Google*, Case Number: 1:23-cv-00108-LMB-JFA, including:

1. Declarations of Google Employees
2. 2023.12.22 Expert Report of Gabriel Weintraub, GOOG-AT-MDL-C-000018734
3. 2023.12.22 Expert Report of R. Ravi, GOOG-AT-MDL-C-000019017
4. 2023.12.22 Expert Report of Robin S. Lee, GOOG-AT-MDL-C-000019273
5. 2023.12.22 Expert Report of Rosa Abrantes-Metz, GOOG-AT-MDL-C-000019786
6. 2023.12.22 Expert Report of Thomas S. Respess, GOOG-AT-MDL-C-000020106
7. 2023.12.22 Expert Report of Timothy Simcoe, GOOG-AT-MDL-C-000020274
8. 2024.01.13 Errata to Abrantes-Metz Expert Report, GOOG-AT-MDL-C-000020435
9. 2024.01.13 Errata to Ravi Expert Report, GOOG-AT-MDL-C-000020437
10. 2024.01.13 Errata to Respess Expert Report, GOOG-AT-MDL-C-000020440
11. 2024.01.13 Errata to Simcoe Expert Report, GOOG-AT-MDL-C-000020467
12. 2024.01.13 Errata to Weintraub Expert Report, GOOG-AT-MDL-C-000020471
13. 2024.01.23 Chevalier Expert Report, GOOG-AT-MDL-C-000020474
14. 2024.01.23 Ferrante Expert Report, GOOG-AT-MDL-C-000020714
15. 2024.01.23 Ghose Expert Report, GOOG-AT-MDL-C-000020767
16. 2024.01.23 Israel Expert Report, GOOG-AT-MDL-C-000021036
17. 2024.01.23 Milgrom Expert Report, GOOG-AT-MDL-C-000021794
18. 2024.01.23 Rinard Expert Report, GOOG-AT-MDL-C-000022191
19. 2024.01.23 Shirky Expert Report, GOOG-AT-MDL-C-000022229
20. 2024.01.23 Simonson Expert Report, GOOG-AT-MDL-C-000022290
21. 2024.01.23 Skinner Expert Report, GOOG-AT-MDL-C-000022948
22. 2024.02.13 Expert Rebuttal Report of Adoria Lim, GOOG-AT-MDL-C-000023002
23. 2024.02.13 Expert Rebuttal Report of Gabriel Weintraub, GOOG-AT-MDL-C-000023226
24. 2024.02.13 Expert Rebuttal Report of Kenneth Wilbur, GOOG-AT-MDL-C-000023322
25. 2024.02.13 Expert Rebuttal Report of R. Ravi, GOOG-AT-MDL-C-000023435
26. 2024.02.13 Expert Rebuttal Report of Robin S. Lee, GOOG-AT-MDL-C-000023516
27. 2024.02.13 Expert Rebuttal Report of Rosa Abrantes-Metz, GOOG-AT-MDL-C-000023887

28. 2024.02.13 Expert Rebuttal Report of Timothy Simcoe, GOOG-AT-MDL-C-000024064
29. 2024.02.13 Expert Rebuttal Report of Wayne Hoyer, GOOG-AT-MDL-C-000024138
30. 2024.02.13 Expert Rebuttal Report of Wenke Lee, GOOG-AT-MDL-C-000024270
31. 2024.02.16 Errata to Ravi Rebuttal Report, GOOG-AT-MDL-C-000024387
32. 2024.02.20 Errata to Simcoe Rebuttal Report, GOOG-AT-MDL-C-000024389
33. 2024.02.23 Errata to Weintraub Rebuttal Report, GOOG-AT-MDL-C-000024390
34. 2024.02.23 Supplemental Errata to Weintraub Expert Report, GOOG-AT-MDL-C-000024391
35. 2024.02.24 Errata to Wilbur Rebuttal Report, GOOG-AT-MDL-C-000024392
36. 2024.02.26 Errata to Hoyer Rebuttal Report, GOOG-AT-MDL-C-000024397
37. 2024.02.28 Errata to Abrantes-Metz Rebuttal Report, GOOG-AT-MDL-C-000024399
38. 2024.03.04 Expert Supplemental Report of Robin S. Lee, GOOG-AT-MDL-C-000024403
39. 2024.03.08 Consolidated Errata to Lee Rebuttal Report, GOOG-AT-MDL-C-000024436
40. 2024.01.13 Expert Report of Weintraub Errata, GOOG-AT-MDL-C-000040965
41. 2024.01.13 Expert Report of Simcoe Errata, GOOG-AT-MDL-C-000040961
42. 2024.01.13 Expert Report of Respess Errata_with Figure Errata_Redacted, GOOG-AT-MDL-C-000040934
43. 2024.01.13 Expert Report of R Ravi Errata, GOOG-AT-MDL-C-000040931
44. 2024.01.13 Expert Report of Abrantes-Metz Errata, GOOG-AT-MDL-C-000040929
45. 2024.03.08 Consolidated Errata to Lee Rebuttal Report, GOOG-AT-MDL-C-000040926
46. 2024.03.04 Expert Supplemental Report of Robin S. Lee, PhD, GOOG-AT-MDL-C-000040893
47. 2024.02.28 Rebuttal Report Errata of Rosa Abrantes-Metz Signed, GOOG-AT-MDL-C-000040889
48. 2024.02.25 Expert Rebuttal Report of Hoyer Errata, GOOG-AT-MDL-C-000040887
49. 2024.02.24 Wilbur Rebuttal Errata, GOOG-AT-MDL-C-000040882
50. 2024.02.23 Weintraub Rebuttal Report Errata, GOOG-AT-MDL-C-000040881
51. 2024.02.23 Expert Report of Weintraub Supplemental Errata, GOOG-AT-MDL-C-000040880
52. 2024.02.20 Errata to Simcoe Rebuttal Report, GOOG-AT-MDL-C-000040879
53. 2024.02.16 Errata to Ravi Rebuttal Report (Highly Confidential), GOOG-AT-MDL-C-000040877
54. 2024.02.13 Rebuttal Report of Rosa Abrantes-Metz, GOOG-AT-MDL-C-000040700
55. 2024.02.13 Expert Report of Wenke Lee, GOOG-AT-MDL-C-000040583
56. 2024.02.13 Expert Rebuttal Report of Wayne Hoyer, GOOG-AT-MDL-C-000040451
57. 2024.02.13 Expert Rebuttal Report of Timothy Simcoe_Redacted, GOOG-AT-MDL-C-000040377
58. 2024.02.13 Expert Rebuttal Report of Robin S. Lee_Redacted, GOOG-AT-MDL-C-000040006
59. 2024.02.13 Expert Rebuttal Report of R Ravi, GOOG-AT-MDL-C-000039925
60. 2024.02.13 Expert Rebuttal Report of Kenneth Wilbur_Redacted, GOOG-AT-MDL-C-000039812

61. 2024.02.13 Expert Rebuttal Report of Gabriel Weintraub_Redacted, GOOG-AT-MDL-C-000039716
62. 2024.02.13 Expert Rebuttal Report of Adoria Lim_Redacted, GOOG-AT-MDL-C-000039492
63. 2024.01.23 Expert Report of William Clay Shirky, GOOG-AT-MDL-C-000039431
64. 2024.01.23 Expert Report of Paul R. Milgrom, GOOG-AT-MDL-C-000039034
65. 2024.01.23 Expert Report of Martin C. Rinard, GOOG-AT-MDL-C-000038996
66. 2024.01.23 Expert Report of Mark A. Israel_Redacted, GOOG-AT-MDL-C-000038238
67. 2024.01.23 Expert Report of Judith A. Chevalier_Redacted, GOOG-AT-MDL-C-000037998
68. 2024.01.23 Expert Report of Itamar Simonson, GOOG-AT-MDL-C-000037340
69. 2024.01.23 Expert Report of Douglas Skinner, GOOG-AT-MDL-C-000037286
70. 2024.01.23 Expert Report of Anthony J. Ferrante, GOOG-AT-MDL-C-000037233
71. 2024.01.23 Expert Report of Anindya Ghose_Redacted, GOOG-AT-MDL-C-000036954
72. 2023.12.22 Expert Report of Timothy Simcoe_Redacted, GOOG-AT-MDL-C-000036793
73. 2023.12.22 Expert Report of Thomas Respess_Redacted, GOOG-AT-MDL-C-000036625
74. 2023.12.22 Expert Report of Rosa Abrantes-Metz_Redacted, GOOG-AT-MDL-C-000036305
75. 2023.12.22 Expert Report of Robin S. Lee, PhD_Redacted, GOOG-AT-MDL-C-000035792
76. 2023.12.22 Expert Report of R Ravi_Redacted, GOOG-AT-MDL-C-000035536
77. 2023.12.22 Expert Report of Gabriel Weintraub_Redacted, GOOG-AT-MDL-C-000035253

<u>Bates Stamped Productions</u>, including:

1. DM_GOOG_0027446
2. FBDOJ012641326 / MetaTx_000000680
3. FBDOJGOOG_00478712
4. FBDOJGOOG_00986815
5. FBDOJGOOG_01484801
6. FBTEX_00078430
7. FBTEX_00079937
8. FBTEX_00116640 / FBTEX_00116632
9. FBTEX_00277880
10. FBDOJGOOG_00327692
11. FBTEX_00327637 / FBTEX_00327634
12. FBTEX_00327690
13. FBTEX_00334404
14. FBTEX_00482531
15. FBTEX_00528526
16. FBTEX_00540345
17. FBTEX_00585545 / FBTEX_00585544
18. FBTEX_00797666
19. FBTEX_00808277
20. FBTEX_00813671
21. FBTEX_00837148
22. FBTEX_00854833
23. FBTEX_00892345
24. FBTEX_00969349
25. FBTEX_01003026
26. FBTEX_01021169
27. FBTEX_01062704
28. FBTEX_01064247
29. FBTEX_01064318
30. FBTEX_01080688
31. FBTEX_01082050
32. FBTEX_01089475
33. FBTEX_01103919
34. FBTEX_01131531
35. FBTEX_01200160
36. FBTEX_01254367
37. FBTEX_01274568
38. FBTEX_00528334
39. FTC_US-GOOGLE-000004531
40. GOOG-AT-MDL-001263607
41. GOOG-AT-MDL-001390730
42. GOOG-AT-MDL-001391213
43. GOOG-AT-MDL-001933227
44. GOOG-AT-MDL-002105969
45. GOOG-AT-MDL-002105984
46. GOOG-AT-MDL-002124829
47. GOOG-AT-MDL-002390899
48. GOOG-AT-MDL-002393442
49. GOOG-AT-MDL-003161451

50. GOOG-AT-MDL-004074544
51. GOOG-AT-MDL-004232880
52. GOOG-AT-MDL-004233138
53. GOOG-AT-MDL-004300268
54. GOOG-AT-MDL-004416785
55. GOOG-AT-MDL-004436768
56. GOOG-AT-MDL-004555181
57. GOOG-AT-MDL-006099844
58. GOOG-AT-MDL-006161050
59. GOOG-AT-MDL-006334729
60. GOOG-AT-MDL-006873424
61. GOOG-AT-MDL-006966530
62. GOOG-AT-MDL-007175167
63. GOOG-AT-MDL-007343585
64. GOOG-AT-MDL-007346556
65. GOOG-AT-MDL-007364833
66. GOOG-AT-MDL-007375672
67. GOOG-AT-MDL-007387750
68. GOOG-AT-MDL-007397182
69. GOOG-AT-MDL-007397197
70. GOOG-AT-MDL-008148533 / GOOG-AT-MDL-008148529
71. GOOG-AT-MDL-008517788
72. GOOG-AT-MDL-008588684
73. GOOG-AT-MDL-008682082 / GOOG-AT-MDL-008682071
74. GOOG-AT-MDL-008754374
75. GOOG-AT-MDL-008835346
76. GOOG-AT-MDL-008858602
77. GOOG-AT-MDL-008881206
78. GOOG-AT-MDL-008886980
79. GOOG-AT-MDL-008953893
80. GOOG-AT-MDL-008964888
81. GOOG-AT-MDL-008979664
82. GOOG-AT-MDL-008991390
83. GOOG-AT-MDL-009026140
84. GOOG-AT-MDL-009289718
85. GOOG-AT-MDL-009291120
86. GOOG-AT-MDL-009299907
87. GOOG-AT-MDL-009321580
88. GOOG-AT-MDL-009429957
89. GOOG-AT-MDL-012512067
90. GOOG-AT-MDL-012514705
91. GOOG-AT-MDL-012524006
92. GOOG-AT-MDL-012549335
93. GOOG-AT-MDL-012693796
94. GOOG-AT-MDL-012767138
95. GOOG-AT-MDL-012837016
96. GOOG-AT-MDL-012857198
97. GOOG-AT-MDL-013290688
98. GOOG-AT-MDL-013291089

99. GOOG-AT-MDL-013292974
100. GOOG-AT-MDL-013299524
101. GOOG-AT-MDL-013299531
102. GOOG-AT-MDL-013300202
103. GOOG-AT-MDL-013378392
104. GOOG-AT-MDL-013908958
105. GOOG-AT-MDL-013918668
106. GOOG-AT-MDL-014427012
107. GOOG-AT-MDL-014460206
108. GOOG-AT-MDL-014462378
109. GOOG-AT-MDL-014486274
110. GOOG-AT-MDL-014524447
111. GOOG-AT-MDL-014618288
112. GOOG-AT-MDL-015241235
113. GOOG-AT-MDL-015622194
114. GOOG-AT-MDL-015844174
115. GOOG-AT-MDL-015929587
116. GOOG-AT-MDL-015997353
117. GOOG-AT-MDL-016457027
118. GOOG-AT-MDL-016534880
119. GOOG-AT-MDL-016627159
120. GOOG-AT-MDL-016656237
121. GOOG-AT-MDL-016772599
122. GOOG-AT-MDL-016838311
123. GOOG-AT-MDL-016924839
124. GOOG-AT-MDL-016937590
125. GOOG-AT-MDL-016943922
126. GOOG-AT-MDL-016967094
127. GOOG-AT-MDL-017187837
128. GOOG-AT-MDL-017394050
129. GOOG-AT-MDL-017494582
130. GOOG-AT-MDL-017664768
131. GOOG-AT-MDL-017746412
132. GOOG-AT-MDL-017749638
133. GOOG-AT-MDL-017762649
134. GOOG-AT-MDL-017864022
135. GOOG-AT-MDL-018248228
136. GOOG-AT-MDL-018427318
137. GOOG-AT-MDL-018448707
138. GOOG-AT-MDL-018548592
139. GOOG-AT-MDL-018618351
140. GOOG-AT-MDL-018652651
141. GOOG-AT-MDL-018998910
142. GOOG-AT-MDL-019001498
143. GOOG-AT-MDL-019306356
144. GOOG-AT-MDL-019386250
145. GOOG-AT-MDL-019552139
146. GOOG-AT-MDL-019571201
147. GOOG-AT-MDL-019588187

148. GOOG-AT-MDL-019633443
149. GOOG-AT-MDL-019642313
150. GOOG-AT-MDL-019653406
151. GOOG-AT-MDL-019721340
152. GOOG-AT-MDL-019767203
153. GOOG-AT-MDL-B-000134141
154. GOOG-AT-MDL-B-001084151
155. GOOG-AT-MDL-B-001140202
156. GOOG-AT-MDL-B-002087955
157. GOOG-AT-MDL-B-002088697
158. GOOG-AT-MDL-B-002088752
159. GOOG-AT-MDL-B-002088926
160. GOOG-AT-MDL-B-002090567
161. GOOG-AT-MDL-B-002091565
162. GOOG-AT-MDL-B-002095353
163. GOOG-AT-MDL-B-002095501
164. GOOG-AT-MDL-B-002095769
165. GOOG-AT-MDL-B-002097533
166. GOOG-AT-MDL-B-002097570
167. GOOG-AT-MDL-B-002097648
168. GOOG-AT-MDL-B-002098265
169. GOOG-AT-MDL-B-002099366
170. GOOG-AT-MDL-B-002105135
171. GOOG-AT-MDL-B-002500395
172. GOOG-AT-MDL-B-002514153
173. GOOG-AT-MDL-B-002547489
174. GOOG-AT-MDL-B-002552122
175. GOOG-AT-MDL-B-002624643
176. GOOG-AT-MDL-B-002760309
177. GOOG-AT-MDL-B-002762758
178. GOOG-AT-MDL-B-002763194
179. GOOG-AT-MDL-B-002764178
180. GOOG-AT-MDL-B-002764191
181. GOOG-AT-MDL-B-002797051
182. GOOG-AT-MDL-B-002803919
183. GOOG-AT-MDL-B-002837096
184. GOOG-AT-MDL-B-003181628
185. GOOG-AT-MDL-B-003735823
186. GOOG-AT-MDL-B-003741154
187. GOOG-AT-MDL-B-004008753
188. GOOG-AT-MDL-B-004015880
189. GOOG-AT-MDL-B-004016318
190. GOOG-AT-MDL-B-004243337
191. GOOG-AT-MDL-B-004247242
192. GOOG-AT-MDL-B-004265772
193. GOOG-AT-MDL-B-004425247
194. GOOG-AT-MDL-B-004680051
195. GOOG-AT-MDL-B-005083974
196. GOOG-AT-MDL-B-005167304

197. GOOG-AT-MDL-B-005168118
198. GOOG-AT-MDL-B-005170475
199. GOOG-AT-MDL-B-005282318
200. GOOG-AT-MDL-B-005372599
201. GOOG-AT-MDL-B-005457387
202. GOOG-AT-MDL-B-006069467
203. GOOG-AT-MDL-B-006316352
204. GOOG-AT-MDL-B-006365895
205. GOOG-AT-MDL-B-006365981
206. GOOG-AT-MDL-B-006939056
207. GOOG-AT-MDL-B-007212533
208. GOOG-AT-MDL-B-007229334
209. GOOG-AT-MDL-B-007232867
210. GOOG-AT-MDL-B-007353902
211. GOOG-DOJ-12948968
212. GOOG-DOJ-13897780
213. GOOG-DOJ-13899823
214. GOOG-DOJ-13911836
215. GOOG-DOJ-13930748
216. GOOG-DOJ-13940086
217. GOOG-DOJ-14008698
218. GOOG-DOJ-14034714
219. GOOG-DOJ-14113270
220. GOOG-DOJ-14139857
221. GOOG-DOJ-14155066
222. GOOG-DOJ-14156827
223. GOOG-DOJ-14161619
224. GOOG-DOJ-14161943
225. GOOG-DOJ-14232497
226. GOOG-DOJ-14352302
227. GOOG-DOJ-14352774
228. GOOG-DOJ-14365517
229. GOOG-DOJ-14433486
230. GOOG-DOJ-14433633
231. GOOG-DOJ-14435110
232. GOOG-DOJ-14436029
233. GOOG-DOJ-14453674
234. GOOG-DOJ-14458088
235. GOOG-DOJ-14494204
236. GOOG-DOJ-14544715
237. GOOG-DOJ-14544743
238. GOOG-DOJ-14556699
239. GOOG-DOJ-14712739
240. GOOG-DOJ-14717683
241. GOOG-DOJ-14733660
242. GOOG-DOJ-14735212
243. GOOG-DOJ-14735427
244. GOOG-DOJ-14740782
245. GOOG-DOJ-14743636

246. GOOG-DOJ-14870370
247. GOOG-DOJ-15021111
248. GOOG-DOJ-15022600
249. GOOG-DOJ-15022611
250. GOOG-DOJ-15029681
251. GOOG-DOJ-15041717
252. GOOG-DOJ-15042589
253. GOOG-DOJ-15064786
254. GOOG-DOJ-15071642
255. GOOG-DOJ-15073261
256. GOOG-DOJ-15076754
257. GOOG-DOJ-15140608
258. GOOG-DOJ-15173140
259. GOOG-DOJ-15186471
260. GOOG-DOJ-15278949
261. GOOG-DOJ-15371972
262. GOOG-DOJ-15417170
263. GOOG-DOJ-15426837
264. GOOG-DOJ-15427800
265. GOOG-DOJ-15428154
266. GOOG-DOJ-15432090
267. GOOG-DOJ-15433127
268. GOOG-DOJ-15434544
269. GOOG-DOJ-15445478
270. GOOG-DOJ-15526215
271. GOOG-DOJ-15590044
272. GOOG-DOJ-15595543
273. GOOG-DOJ-15597407
274. GOOG-DOJ-15598151
275. GOOG-DOJ-15601229
276. GOOG-DOJ-15616526
277. GOOG-DOJ-15621021
278. GOOG-DOJ-15628088
279. GOOG-DOJ-15631978
280. GOOG-DOJ-15637110
281. GOOG-DOJ-15675958
282. GOOG-DOJ-15733030
283. GOOG-DOJ-15766965
284. GOOG-DOJ-15789579
285. GOOG-DOJ-16271009
286. GOOG-DOJ-16732920
287. GOOG-DOJ-16976651
288. GOOG-DOJ-17101302
289. GOOG-DOJ-17105055
290. GOOG-DOJ-17581278
291. GOOG-DOJ-19090482
292. GOOG-DOJ-19119112
293. GOOG-DOJ-27760500_DUP_1
294. GOOG-DOJ-27762649

295. GOOG-DOJ-27762691
296. GOOG-DOJ-27796586
297. GOOG-DOJ-27803156
298. GOOG-DOJ-27803505
299. GOOG-DOJ-27803533
300. GOOG-DOJ-27804205
301. GOOG-DOJ-27804876
302. GOOG-DOJ-28385887
303. GOOG-DOJ-28420330
304. GOOG-DOJ-29427368
305. GOOG-DOJ-29453844
306. GOOG-DOJ-30066802
307. GOOG-DOJ-32018179
308. GOOG-DOJ-32018452
309. GOOG-DOJ-32022422
310. GOOG-DOJ-32062262
311. GOOG-DOJ-32262980
312. GOOG-DOJ-32265694
313. GOOG-DOJ-AT-00003565
314. GOOG-DOJ-AT-00087809
315. GOOG-DOJ-AT-00252559
316. GOOG-DOJ-AT-00569953
317. GOOG-DOJ-AT-00575435
318. GOOG-DOJ-AT-00588995
319. GOOG-DOJ-AT-00651561
320. GOOG-DOJ-AT-00663174
321. GOOG-DOJ-AT-00750354
322. GOOG-DOJ-AT-01019411
323. GOOG-DOJ-AT-01027937
324. GOOG-DOJ-AT-01514374
325. GOOG-DOJ-AT-01525829
326. GOOG-DOJ-AT-01682499
327. GOOG-DOJ-AT-01811246
328. GOOG-DOJ-AT-01814428
329. GOOG-DOJ-AT-01814688
330. GOOG-DOJ-AT-01815482
331. GOOG-DOJ-AT-01816686
332. GOOG-DOJ-AT-01848255
333. GOOG-DOJ-AT-01933085
334. GOOG-DOJ-AT-02070545
335. GOOG-DOJ-AT-02122748
336. GOOG-DOJ-AT-02146896
337. GOOG-DOJ-AT-02153612
338. GOOG-DOJ-AT-02275074
339. GOOG-DOJ-AT-02275933
340. GOOG-DOJ-AT-02319393
341. GOOG-DOJ-AT-02482767
342. GOOG-DOJ-AT-02634336
343. GOOG-DOJ-AT-02639830

344. GOOG-NE-01787563
345. GOOG-NE-02557667
346. GOOG-NE-02558055
347. GOOG-NE-03231727
348. GOOG-NE-03236272
349. GOOG-NE-03467508
350. GOOG-NE-03616222
351. GOOG-NE-03627597
352. GOOG-NE-03727939
353. GOOG-NE-03844060
354. GOOG-NE-03869994
355. GOOG-NE-03872605
356. GOOG-NE-04304005
357. GOOG-NE-04689402
358. GOOG-NE-05241137
359. GOOG-NE-05270679
360. GOOG-NE-06113062
361. GOOG-NE-06547825
362. GOOG-NE-06568562
363. GOOG-NE-06568719
364. GOOG-NE-06729717
365. GOOG-NE-06841582
366. GOOG-NE-07841238
367. GOOG-NE-07946731
368. GOOG-NE-08116078
369. GOOG-NE-08112779
370. GOOG-NE-09138394
371. GOOG-NE-09698556
372. GOOG-NE-09711564
373. GOOG-NE-10573952
374. GOOG-NE-10660658
375. GOOG-NE-10660882
376. GOOG-NE-10777158
377. GOOG-NE-10952313
378. GOOG-NE-10976657
379. GOOG-NE-10977717
380. GOOG-NE-10978402
381. GOOG-NE-10985565
382. GOOG-NE-11813912
383. GOOG-NE-11847370
384. GOOG-NE-11877282
385. GOOG-NE-12727978
386. GOOG-NE-12738745
387. GOOG-NE-12742008
388. GOOG-NE-12787759
389. GOOG-NE-12885078
390. GOOG-NE-13004102
391. GOOG-NE-13200100
392. GOOG-NE-13205235

393. GOOG-NE-13205865
394. GOOG-NE-13210272
395. GOOG-NE-13327192
396. GOOG-NE-13340752
397. GOOG-NE-13349343
398. GOOG-NE-13367768
399. GOOG-NE-13369624
400. GOOG-NE-13373170
401. GOOG-NE-13374150
402. GOOG-NE-13379438
403. GOOG-NE-13386334
404. GOOG-NE-13389481
405. GOOG-NE-13390045
406. GOOG-NE-13393115
407. GOOG-NE-13401911
408. GOOG-NE-13405025
409. GOOG-NE-13415537
410. GOOG-NE-13541543
411. GOOG-NE-13569193
412. GOOG-NE-13577462
413. GOOG-NE-13589899
414. GOOG-NE-13624783
415. GOOG-TEX-00039266
416. GOOG-TEX-00079344
417. GOOG-TEX-00089835
418. GOOG-TEX-00090151
419. GOOG-TEX-00090282
420. GOOG-TEX-00090969
421. GOOG-TEX-00092657
422. GOOG-TEX-00105202
423. GOOG-TEX-00105361
424. GOOG-TEX-00106255
425. GOOG-TEX-00109167
426. GOOG-TEX-00110068
427. GOOG-TEX-00110540
428. GOOG-TEX-00110544
429. GOOG-TEX-00111494
430. GOOG-TEX-00116069
431. GOOG-TEX-00116639
432. GOOG-TEX-00119737
433. GOOG-TEX-00119845
434. GOOG-TEX-00120626
435. GOOG-TEX-00120775
436. GOOG-TEX-00156142
437. GOOG-TEX-00167119
438. GOOG-TEX-00177559
439. GOOG-TEX-00206687
440. GOOG-TEX-00216163
441. GOOG-TEX-00234150

442. GOOG-TEX-00240572
443. GOOG-TEX-00270127
444. GOOG-TEX-00271177
445. GOOG-TEX-00309326
446. GOOG-TEX-00326649
447. GOOG-TEX-00336527
448. GOOG-TEX-00344083
449. GOOG-TEX-00370306
450. GOOG-TEX-00374779
451. GOOG-TEX-00375239
452. GOOG-TEX-00452866
453. GOOG-TEX-00513684
454. GOOG-TEX-00643890
455. GOOG-TEX-00689539
456. GOOG-TEX-00705131
457. GOOG-TEX-00715805
458. GOOG-TEX-00716782
459. GOOG-TEX-00777573
460. GOOG-TEX-00778301
461. GOOG-TEX-00797340
462. GOOG-TEX-00806640
463. GOOG-TEX-00806682
464. GOOG-TEX-00814407
465. GOOG-TEX-00825713
466. GOOG-TEX-00828547
467. GOOG-TEX-00831660
468. GOOG-TEX-00850729
469. GOOG-TEX-00858576
470. GOOG-TEX-00905673
471. GOOG-TEX-00959457
472. GOOG-TEX-00959461
473. GOOG-TEX-00969525
474. GOOG-TEX-00978814
475. GOOG-TEX-01004466
476. GOOG-TEX-01036150
477. GOOG-TEX-01142635
478. GOOG-TEX-01244428
479. GOOG-TEX-01279945
480. GOOG-TEX-00974499
481. GOOG-TEX-00001418
482. GOOG-TEX-00597317
483. METATX_000000680
484. NEXSTAR090311

APPENDIX C. RELEVANT CONCEPTS IN AUCTION THEORY

1. **Numerical Example for Reserve Prices**. The leading example can be modified and utilized to understand reserve prices as well. Imagine that there are five bidders, who submit bids of $1, $8, $3, $5, $2, respectively. If the auctioneer sets a reserve of $4, then in the first-price auction with reserve $4, Bidder Two wins and pays $8. In the second-price auction with reserve $4, Bidder Two wins and pays $5. That is, both auctions conclude exactly as if there were no reserve, because the reserve is smaller than the second-highest bid. Observe also that Bidder Two's minimum bid to win is $5, and all other bidders have a minimum bid to win of $8. Observe again that Bidder Two's bid ($8) exceeds their minimum bid to win ($5), and therefore Bidder Two wins, whereas all other bidders have bids ($1, $3, $5, $2) below their minimum bid to win ($8) and therefore lose.[1] If the auctioneer sets a reserve $r between $5 and $8, then in the first-price auction with reserve $r, Bidder Two wins and pays $8. That is, the first-price auction concludes exactly as if there were no reserve, because the reserve is smaller than the highest bid. In the second-price Auction with reserve $r, Bidder Two wins and pays $r. That is, Bidder Two still wins the second-price auction, because they outbid the reserve. However, Bidder Two pays more because the reserve $r is treated as the second-highest bid.[2] Observe also that Bidder Two's minimum bid to win is $r, and all other bidders have a minimum bid to win of $8. Observe again that Bidder Two's bid ($8) exceeds their minimum bid to win ($r) and therefore wins, whereas all other bidders have bids ($1, $3, $5, $2) below their minimum bid to win ($8) and therefore lose.[3] Lastly, if the auctioneer sets a reserve $r that is greater than $8, then in both auctions the item remains unsold. Observe also that all bidders have a minimum bid to win of $r and that all bidders have bids ($1, $8, $3, $5, $2) less than their minimum bid to win ($r), and therefore they all lose.[4]

---

[1] Revisiting the first interpretation, Bidder Zero submits a bid of $5. This means that Bidder Two is the highest bidder and Bidder Four is the highest other bidder and therefore Bidder Zero (and the reserve) is irrelevant. Revisiting the second interpretation, Bidders Two and Four both survive the reserve. Therefore, Bidder Two will win the auction (as the highest surviving bidder), and Bidder Four will be the highest other surviving bidder. Therefore, Bidder Two will pay $5 (Bidder Four's bid).

[2] In the event of a tie (in this example, r = $8), it is common to break ties in favor of the bidder, so that a sale occurs.

[3] Revisiting the first interpretation, Bidder Zero submits a bid of r between $5 and $8. This means that Bidder Two is the highest bidder and Bidder Zero is the highest other bidder. Therefore, Bidder Two wins both auctions, pays their bid ($8) in the first-price auction, and pays Bidder Zero's bid ($*r*) in the second-price auction. Revisiting the second interpretation, only Bidder Two survives the reserve. Therefore, Bidder Two will win both auctions. In the first-price auction, Bidder Two pays their bid ($8). In the second- price auction, Bidder Two pays a tentative price of $0 (because there are no other surviving bids, so the highest other surviving bid is $0), which is increased to the price floor of r.

[4] Revisiting the first interpretation, Bidder Zero is the highest bidder, so Bidder Zero wins. This means the item stays unsold and no payments occur. Revisiting the second interpretation, no bidders survive the reserve. Therefore, there are no remaining bidders participating in either the first- or second-price auction, so the item stays unsold, and no payments occur.

2. **Theorem 1** (Truthfulness of Second-Price Auctions)**.** With private values, the second-price auction is truthful. Moreover, any single-item auction with a monotone allocation rule that charges the winner their minimum bid to win is truthful with private values. Therefore, the second-price auction with reserve is also truthful, as is the second-price auction with personalized reserves.

3. *Proof.* Consider any single-item auction with monotone allocation rule that charges the winner their minimum bid to win, and consider Bidder *i*. We know that, given the bids of other bidders, there are two possible outcomes: (a) Bidder *i* might win and pay their minimum bid to win (if they submit a bid above their minimum bid to win), or (b) Bidder *i* might lose (if they submit a bid below their minimum bid to win). If Bidder *i*'s value exceeds their minimum bid to win, then they strictly prefer (a) over (b). Moreover, Bidder *i* can guarantee outcome (a) by submitting any bid above their minimum bid to win. Bidder *i*'s value is one such bid, and therefore we conclude that whenever Bidder i's value exceeds their minimum bid to win, bidding their value achieves the best possible outcome. Similarly, if Bidder *i*'s value falls below their minimum bid to win, then they strictly prefer (b) over (a). Moreover, Bidder *i* can guarantee outcome (b) by submitting any bid below their minimum bid to win. Bidder *i*'s value is one such bid, and therefore we conclude that whenever Bidder i's value falls below their minimum bid to win, bidding their value achieves the best possible outcome. Finally, we conclude that if Bidder *i*'s value is equal to their minimum bid to win, they are indifferent between (a) and (b), and so no matter how ties are broken whenever Bidder i's value is equal to their minimum bid to win, bidding their value achieves the best possible outcome (because either of the two possible outcomes are equally the best possible outcome).

4. In sum, we observe that no matter how the other bidders' bids are fixed, and no matter Bidder *i*'s value, Bidder *i*'s value is either greater than, less than, or equal to the minimum bid to win induced by the other bidders' bids, and in all three cases Bidder *i* achieves the best possible outcome by bidding their value. Therefore, any single-item auction that charges the winner their minimum bid to win is truthful. Importantly, we used the fact that the auction charges the winner their minimum bid to win to conclude that the only two possible outcomes are (a) or (b). If the auction charges the winner alternate prices (such as the first-price auction, which charges the winner their bid), then there are more possible outcomes, and the proof fails.

APPENDIX D. CONDUCT ANALYSIS: DYNAMIC ALLOCATION AND ENHANCED DYNAMIC ALLOCATION

5. Here, I rigorously work through an example where two bidders participate in a first-price auction, but one bidder learns the others' bids before bidding. This appendix is intended for a reader who is familiar with basic auction theory, and in particular understands the concept of Bayes-Nash equilibrium.

6. Consider an auction in the independent private values model where there are 2 bidders whose values are each drawn independently and uniformly from [0,10].

7. **Lemma 1.** With two bidders whose values are each drawn independently and uniformly from [0,10], it is a Bayes-Nash equilibrium of the first-price auction for each bidder to submit a bid of v/2 when their value is v.

8. *Proof.* If one bidder is using this strategy, then the remaining bidder's best response is surely a bid in [0,5] – this is because the bid 5 wins with probability 1, so a higher bid is strictly worse. Now, observe that a bid of b wins if and only if the other bidder has a value of at most $2b$, which occurs with probability $2b/10 = b/5$. Therefore, the remaining bidder's optimal bid maximizes $(v - b) * (b/5)$ over [0,5] (the utility they gain when they win, times the probability that they win). This is optimized at $b = v/2$.[5]

9. **Lemma 2.** With two bidders whose values are each drawn independently and uniformly from [0,10], if Bidder One sees the bid of Bidder Two before submitting their own in a first-price auction, it is a Bayes-Nash equilibrium for Bidder Two to submit a bid of v/2 when their value is v, and for Bidder One to outbid Bidder Two by the minimal amount possible when their value exceeds Bidder Two's bid (and to submit a losing bid otherwise).

10. *Proof.* No matter what Bidder Two does, it is Bidder One's optimal strategy to outbid Bidder Two by the minimal amount possible when their value exceeds Bidder Two's bid (and to submit a losing bid otherwise). Therefore, we just need to confirm that Bidder Two is best responding. Bidder Two's bid of b wins if and only if Bidder One's value is at most b, which occurs with

[5] To see this, take the derivative with respect to b, which is $-b/5+(v-b)/5 = v/5-2b/5$. This is strictly positive for $b<v/2$, and strictly negative for $b>v/2$, and therefore the initial function is optimized at $b = v/2$.

probability $b/10$. Therefore, Bidder Two's optimal bid maximizes $(v - b) * b/10$ over [0,10]. This is optimized at $b = v/2$.[6]

11. **Corollary 1.** With two bidders whose values are each drawn independently and uniformly from [0,10], in both the sealed bid first-price auction and the first-price auction where Bidder One sees Bidder Two's bid, Bidder Two submits a bid of v/2 when their value is v in Bayes-Nash equilibrium. The only distinction is that in a sealed bid first-price auction, Bidder One also submits a bid of v/2 when their value is v, whereas Bidder One submits the minimum bid to win if they can see Bidder Two's bid (if their value exceeds Bidder Two's bid – otherwise they submit a losing bid).

12. Now that we have identified a Bayes-Nash equilibrium in both auction formats, we can compare them. Note that it is not crucial for the example that Bidder Two happens to behave the same way in both Bayes-Nash equilibria (this is a coincidence, rather than a key feature) – what is important is only that we have found both Bayes-Nash equilibria.

13. **Lemma 3.** With two bidders whose values are each drawn independently and uniformly from [0,10], the sealed bid first-price auction achieves an expected revenue of 10/3 ≈ 3.33 in Bayes-Nash and the first-price auction where Bidder One sees Bidder Two's bid achieves an expected revenue of 2.5 in Bayes-Nash.

14. *Proof.* The winning bid in the sealed bid first-price auction in Bayes-Nash is the maximum of the two shaded bids. Each shaded bid is drawn independently and uniformly from [0,5], and the expected maximum of two independent, uniform draws from [0,5] is $10/3$. The winning bid in the first-price auction where Bidder One sees Bidder Two's bid is Bidder Two's bid (either because Bidder Two wins, or because Bidder One marginally outbid them). Bidder Two's bid is drawn uniformly from [0,5], and the expected value of their bid is therefore 2.5.

15. **Lemma 4.** With two bidders whose values are each drawn independently and uniformly from [0,10], each bidder wins half the time in Bayes-Nash of the first-price auction, and Bidder One wins 3/4 of the time in Bayes-Nash of the first-price auction where Bidder One sees Bidder

[6] To see this, take the derivative with respect to b, which is $-b/10+(v-b)/10 = v/10-b/5$. This is strictly positive for $b<v/2$, and strictly negative for $b>v/2$, and therefore the initial function is optimized at $b = v/2$.

Two's bid. In particular, Bidder One wins all of the "high-value" auctions for which they have a value of at least 5, and half of the "low-value" auctions for which they have a value of at most 5.

16. *Proof.* The sealed bid first-price auction is symmetric, and so is the Bayes-Nash, so each bidder wins exactly half the time. Bidder Two's bid in the first-price auction where Bidder One sees Bidder Two's bid is drawn independently and uniformly from [0,5], whereas Bidder One's value is drawn independently and uniformly from [0,10]. Therefore, if Bidder One's value is ≥ 5, they certainly win, and this happens with probability 1/2. If Bidder One's value is ≤ 5, then their value is drawn uniformly from [0,5] conditioned on this, and therefore both Bidder One's value and Bidder Two's bid are drawn independently and uniformly from [0,5]. Therefore, the maximum is uniformly at random, and Bidder One wins half the time. Putting both together, Bidder One wins all the time when their value exceeds 5 (which happens with probability 1/2), and half the time when their value is below 5 (which happens with probability 1/2), and therefore the total probability that Bidder One wins is 3/4.

17. **Lemma 5.** With two bidders whose values are each drawn independently and uniformly from [0,10], the expected welfare[7] in Bayes-Nash equilibrium of the first-price auction is 20/3 ≈ 6.67, and the expected welfare in Bayes-Nash equilibrium of the first-price auction where Bidder One sees Bidder Two's bid is 6.25.

18. *Proof.* In Bayes-Nash equilibrium of the sealed bid first-price auction, the highest-value bidder wins. The expected highest value is the expected maximum of two independent, uniform draws from [0,10], which is 20/3. In Bayes-Nash equilibrium of the first-price auction where Bidder One sees Bidder Two's bid, Bidder One wins whenever their value exceeds 5. This happens with probability 1/2, and their expected value conditioned on this is 7.5. When Bidder One's value falls below 5, each bidder is equally likely to win, and the expected highest bid is the maximum of two independent uniform draws from [0,5]. Half of the time, independent of what the highest bid is, this bid will come from Bidder One and be equal to their value. The other half of the time, this bid will come from Bidder Two and be equal to half their value. Therefore, half of the time the winning bid is equal to the winning bidder's value, and the other half of the time it is equal to half their value. This means that the winning bidder's value is in expectation equal to 3/2 times the winning bid. The winning bid in this case is 10/3 in expectation, and therefore the expected winning

[7] "Welfare" refers to the value of the bidder who wins the item (and is 0 in case no bidder wins the item).

bidder's value in this case is 10/3 * 3/2 = 5. Putting both cases together, this means that the expected value of the winning bidder is 7.5 * 1/2+5 * 1/2 = 6.25.

19. **Lemma 6.** With two bidders whose values are each drawn independently and uniformly from [0,10], each bidder has expected value of 10/3 ≈ 3.33 and expected utility of 5/3 ≈ 1.67 in Bayes-Nash of the first-price auction. In Bayes-Nash Equilibrium of the first-price auction where Bidder One sees Bidder Two's bid, Bidder One has expected value of 55/12 ≈ 4.58 and expected utility of 35/12 ≈ 2.92, and Bidder Two has expected value of 5/3 ≈ 1.67 and expected utility of 5/6 ≈ .83.

20. *Proof.* The value claim for the sealed bid first-price auction follows as the expected welfare is 20/3, which must be split evenly among the two bidders due to symmetry for 10/3 each. The utility claim further follows as the expected revenue is 10/3, which must also be split evenly among the two bidders due to symmetry, leaving a utility of 5/3 for each. For the first-price auction where Bidder One sees Bidder Two's bid,to compute the expected value of Bidder One, we first observe that they win with probability one whenever their value exceeds 5. This happens half the time, and their expected value conditioned on this is 7.5. We next observe that when their value is ≤ 5, both their value and Bidder Two's bid are drawn independently and uniformly from [0,5], and the highest wins. The expected maximum is 10/3 and is equal to Bidder One half the time. Therefore, Bidder One's expected value from this case is 5/3. Summing both cases together, we get that Bidder One's expected value from participating is 7.5 * 1/2 + 5/3 * 1/2 = 55/12 ≈ 4.58. To compute the expected value of Bidder Two, we first observe that they can only win when Bidder One's value is less than 5, which happens half the time. Conditioned on this, the expected winning bid is the expected maximum of two independent, uniform draws from [0,5], which is 10/3, and equal to Bidder Two half the time. When Bidder Two wins, their bid is equal to half their value, meaning that Bidder Two's expected value, conditioned on winning, is 20/3, and Bidder Two wins 1/4 of the time. This means their expected value is 5/3 ≈ 1.67. To compute the expected payment of both bidders, recall that the payment is always Bidder Two's bid, and Bidder Two makes that payment if and only if Bidder One's value is lower. Bidder Two's bid is uniformly distributed on [0,5], and the probability that Bidder One's value is lower than Bidder Two's bid of *b* is *b*/10. Therefore, the expected payment made by Bidder Two is exactly $\int_0^5 b * \left(\frac{b}{10}\right) * \left(\frac{1}{5}\right) db = b^3/150|_0^5 = 5/6 \approx .83$ (integrate the probability that Bidder Two's bid is *b*, which yields 1/5 *db*, times the payment when the bid is *b*, *b*, times the probability that Bidder Two pays the bid

conditioned on this, $b/10$). The expected payment of Bidder One is then equal to the total revenue (2.5) minus Bidder Two's expected payment (5/6), which is $5/3 \approx 1.67$. Once we have the expected payment of each bidder, their respective utilities are just the difference between their expected value and expected payment – $55/12 - 5/3 = 35/12$ and $5/3 - 5/6 = 5/6$.

APPENDIX E. CONDUCT ANALYSIS: EXCHANGE BIDDING

21. In this appendix, I explain how header bidding with (Enhanced) Dynamic Allocation generates more revenue for publishers than waterfalling with (Enhanced) Dynamic Allocation.[8]

22. First, observe that any non-Google exchange will again naturally submit a higher bid under header bidding than under the waterfall process with (Enhanced) Dynamic Allocation. To see this, observe first that non-Google exchanges are always called under header bidding (whereas they may not be called under waterfalling). Next, observe that a non-Google exchange called under waterfalling need only exceed its reserve in order to win, whereas a non-Google exchange called under header bidding must exceed not only its reserve but also all other non-Google exchanges (*and* put pressure against AdX's Last Look).

23. Moreover, on the same bids it's again the case that the winning payment under header bidding with (Enhanced) Dynamic Allocation exceeds that under waterfalling with (Enhanced) Dynamic Allocation. This will follow by considering a few cases.

24. Imagine that the maximum bid that clears its personalized reserve is h, AdX's reserve (which already takes a maximum with the highest temporary CPM) is r, and the maximum value in AdX is v. Recall that the reserve to AdX under header bidding with (Enhanced) Dynamic Allocation is the maximum of h and r, and the reserve to AdX under waterfalling with (Enhanced) Dynamic Allocation is r. There are a few cases to consider.

25. First, perhaps v exceeds the maximum of h and r. In this case, AdX certainly wins under both auction formats, and pays more under header bidding with (Enhanced) Dynamic Allocation because AdX's reserve is higher than under waterfalling with (Enhanced) Dynamic Allocation.

26. Next, perhaps h exceeds the maximum of v and r. In this case, header bidding with (Enhanced) Dynamic Allocation achieves revenue exactly h, while waterfalling with (Enhanced) Dynamic Allocation achieves revenue no more than h (clearly AdX will not bid more than h, and any non-Google exchange also gives payment no more than h).

[8] Again, I consider the case where the publisher sets identical reserves under both the waterfall and header bidding. A publisher who further optimizes reserves under header bidding could only further increase their revenue. I also consider a publisher who uses default Value CPMs for header bids. An optimizing publisher could again only further increase their revenue.

27. Finally, perhaps r exceeds the maximum of h and v. In this case, AdX certainly does not win in either auction format, because its highest value is less than its reserve. There are two subcases to consider here, depending on the highest temporary CPM from high priority line items, t. If h exceeds t, then header bidding with (Enhanced) Dynamic Allocation achieves a revenue of exactly h in this case, whereas waterfalling with (Enhanced) Dynamic Allocation achieves a revenue of at most h (because the winning exchange is *some* exchange that exceeds its reserve, or perhaps the high priority line item wins before visiting an exchange that exceeds its reserve, which gives $t < h$). Or perhaps t exceeds h. In this case, the high priority line item wins in both formats and the revenues are equal.

28. Taking all of the above cases together, we have now shown that with the same non-Google bids, Header Bidding with (Enhanced) Dynamic Allocation generates greater publisher payout than waterfalling with (Enhanced) Dynamic Allocation. Our first paragraph also establishes that non-Google bids under Header Bidding with (Enhanced) Dynamic Allocation will also be at least as high as non-Google bids under waterfalling with (Enhanced) Dynamic Allocation. Therefore, Header Bidding with (Enhanced) Dynamic Allocation generates more revenue than waterfalling with (Enhanced) Dynamic Allocation.

APPENDIX F. CONDUCT ANALYSIS: UNIFIED PRICING RULES

29. This appendix provides proofs of claims made in the body. Aspects of this section are intended to be accessible for a representative reader, while some aspects are intended for a reader with background in auction theory (and in particular, who is familiar with the concept of 'revenue curves' following Myerson (1981)).[9]

1) Analyzing reserves under treat-as-single and eCPM heuristics

30. There are n bidders, and the auctioneer has a prior D_i over the bids they expect to receive from Bidder *i* (and the distributions are independent). The auctioneer is running a first-price auction with personalized or uniform reserves and uses some heuristic to set the reserves (because optimizing reserves is hugely complex, especially due to complications caused by strategic bidding in first-price auctions). Importantly, the "market primitives" $D_1,\ldots,D_n$ are the same, and the only difference is how reserves are set.

31. **eCPM Heuristic:** Under the **eCPM heuristic**, the auctioneer sets reserves based on data from past auctions. The personalized reserve for Bidder *i* is equal to the average payment that Bidder *i* made while winning similar impressions in the past, plus the average harm caused due to ad quality by the ads displayed by Bidder *i* when while winning similar impressions in the past. The auctioneer sets unified reserves equal to the average payment made for similar impressions in the past, plus the average harm caused due to ad quality by the ads displayed while winning similar impressions in the past. To be extra clear, if f_i denotes the fraction of past similar impressions that Bidder i won, r_i denotes the average payment made by Bidder i while winning these ads, and h_i denotes the average harm caused due to ad quality by Bidder *i*'s ads, then under the eCPM heuristic the personalized reserve for bidder *i* is r_i+h_i, and the unified reserve is $\sum_{i=1}^{n} f_i * (r_i + h_i)$.

32. **Treat-as-Single Heuristic:** Under the **treat-as-single heuristic**, the auctioneer optimally solves a related single-bidder optimization (even though this is unlikely to solve the joint optimization – it is a commonly-studied heuristic. See (Hartline and Roughgarden 2009)[10] and subsequent works). For each distribution D_i, the auctioneer computes the optimal reserve r_i, which

[9] Roger B. Myerson. "Optimal Auction Design." *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73. Subsequent works are also relevant.
[10] Hartline and Roughgarden. "Simple versus Optimal Mechanisms." *Proceedings of the 10th ACM Conference on Electronic Commerce*. 2009. pg. 225-234.

optimizes r*Pr[bid drawn from D_i exceeds r] over all possible reserves r. This optimizes the auctioneer's revenue in the case that Bidder i is the only bidder and is used in seminal work of Myerson (1981)[11] and subsequent works. The treat-as-single heuristic sets personalized reserves by simply using r_i for each Bidder i. For its unified reserve, the treat-as-single heuristic first imagines a hypothetical "average bidder" that produces bids distributed according to the average of the D_is, D. That is, a random draw from the distribution D first picks a uniformly random bidder i and then draws from D_i. Equivalently, the CDF of D is equal to the average of the CDFs of each D_i.[12] The treat-as-single heuristic sets its unified reserve equal to the revenue-optimal reserve for D.

33. **Proposition:** For all instances, the unified reserve set by the eCPM heuristic lies between its minimum and maximum personalized reserves. In particular, whatever bidder faces the highest personalized reserve faces a lower unified reserve, and whatever bidder faces the lowest personalized reserve faces a higher unified reserve.

34. **Proof:** The proposition follows immediately from the formula for the unified reserve. Observe that the formula for the unified reserve is a weighted average of the personalized reserves. A weighted average must lie somewhere between the minimum and maximum.

35. **Proposition:** If each D_i is *regular*,[13] then the unified reserve set by the treat-as-single heuristic lies between its minimum and maximum personalized reserves. In particular, whatever bidder faces the highest personalized reserve faces a lower unified reserve, and whatever bidder faces the lowest personalized reserve faces a higher unified reserve.

36. **Proof:** Let me first observe a few facts regarding revenue-optimal reserves.

a. First, the revenue-optimal reserve for a distribution Di can be found in the following manner using its revenue curve. Find the probability of sale q that maximizes $q*F^{-1}(1-q)$, and then set reserve $F^{-1}(q)$. This follows immediately from the definition of the revenue curve, and the definition of the revenue-optimal Myerson reserve.

---

[11] Roger B. Myerson. "Optimal Auction Design." *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73. Subsequent works are also relevant.

[12] Note that the subsequent claim I make, and its proof, hold verbatim if instead I let D be a weighted average of the D_is. for simplicity of notation, I focus on the unweighted average.

[13] A distribution with CDF F is regular if its revenue curve $q*F^{-1}(1-q)$ is concave on [0,1]. Regularity is commonly assumed on distributions in auction theory. *See* Roger B. Myerson. "Optimal Auction Design." *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73. Subsequent works are also relevant.

b. Second, if D_i is regular, then the revenue function is non-increasing above the optimal reserve, and non-decreasing below it. That is, if we start from a reserve of zero, and slowly increase our reserve until we hit the optimum, our revenue (assuming D_i is the only bidder) will weakly increase, if D_i is regular. This follows by concavity of the revenue curve (and is not guaranteed to hold otherwise). Similarly, once we pass the optimum and continue increasing our reserve, our revenue (assuming D_i is the only bidder) will weakly decrease. That is, if we consider the function $REV_i(p)$ to denote p*Pr[bid drawn from Di exceeds p], then $REV_i(p)$ is weakly increasing on $[0,r_i]$ and weakly decreasing on $[r_i,\infty)$ for any regular distribution D_i, where r_i denotes the optimum of REV_i (and the optimal reserve).

c. Now, let us consider setting a reserve on D. Observe that the revenue function REV(p) for D is equal to p*Pr[bid drawn from D exceeds p]= $p^*\sum_{i=1}^{n} Pr[bid\ drawn\ from\ D_i\ exceeds\ p]/n\ = \sum_{i=1}^{n} REV_i(p)/n$. Therefore, in any range where each revenue function REV_i is non-decreasing, REV is non-decreasing as well, and in any range where each REV_i is non-increasing, REV is non-increasing as well. In particular, because each REV_i is non-decreasing below the minimum personalized reserve, REV is non-decreasing as well (which implies that the optimum is at least as large as the minimum personalized reserve). Similarly, because each REV_i is non-increasing above the maximum personalized reserve, REV is non-increasing as well (which implies that the optimum is no larger than the maximum personalized reserve).

d. Therefore, we conclude that when each D_i is regular, the treat-as-single heuristic sets a unified reserve somewhere between its minimum and maximum personalized reserve.[14]

[14] Note that the only property required of each D_i for our proof is that the revenue function be non-decreasing below the optimum and non-increasing above the optimum. Regularity is a sufficient condition, but other conditions suffice as well, such as declining marginal revenues (Laffont and Robert, 1996). *See* Jean-Jacques Laffont and Jacques Robert. "Optimal auction with financially constrained buyers." *Economics Letters* vol. 52, no. 2. 1996. pg. 181-186. In particular, our conclusions hold more broadly than when each prior is regular, but we state our conclusions for regular distributions because it is the most well-studied condition in auction theory.

2) On efficiency outcomes

37. Equilibria of first-price auctions are notoriously complex, and it is not guaranteed that the bidder with the highest value will win the item. Therefore, no purely mathematical analysis can confidently predict exactly how bidders will bid. Jin and Lu (2023) establish that there is always a fully efficient *Bayes-correlated equilibrium* in first-price auctions. Therefore, I do not claim to predict that real-life bidding behavior in first-price auctions results in efficient outcomes, but only posit that it is a natural assumption for mathematical investigation.

3) Impact of personalized versus unified reserves on bidders

38. Below, I establish claims cited in the body of this section of the report concerning the impact of personalized versus unified reserves on the utility of a bidder.

39. **Proposition:** Consider a list of n bidders who face personalized reserves $r_1,\ldots,r_n$, and let r_i denote the maximum reserve. Let also $r \leq r_i$ be a unified reserve. For any values $v_1,\ldots,v_n$, if Bidder i wins the item in a second-price auction with personalized reserves $r_1,\ldots,r_n$ and pays p, then Bidder i also wins the item in a second-price auction with unified reserve r and pays at most p.

40. **Proof:** The winner of the second-price auction with personalized reserves $r_1,\ldots,r_n$ is the highest bidder who exceeds their reserve, and the payment is their minimum bid to win. If this is Bidder i, this means that v_i exceeds r_i, and no other bid exceeds v_i (because any bid of Bidder j that exceeds v_i certainly exceeds r_j, as $r_{j \leq} r_i$). If no other bidder exceeds r_i, then Bidder i's minimum bid to win is r_i. If the highest other bidder exceeds r_i, then this bidder certainly exceeds their own reserve, and therefore the highest other bid becomes Bidder i's minimum bid to win.

41. Let us now consider the same values in a second-price auction with unified reserve r. Because $r \leq r_i$, and Bidder i is the highest bidder, and $v_i \geq r_i$, this means that Bidder i still wins under unified reserves, as desired. Moreover, the price paid by Bidder i is now equal to r, if no other bid exceeds r, and the highest other bid otherwise. In particular, if the highest other bid exceeds both r_i and r, then Bidder i pays the same under both personalized and unified reserves. If the highest other bid exceeds r but not r_i, Bidder i pays more under personalized reserves (r_i) than unified reserves (highest other bid). If the highest other bid does not exceed r, Bidder i again pays more under personalized reserves (r_i) than unified reserves (r).

42. Therefore, we conclude that whenever Bidder i would win under personalized reserves, Bidder i certainly wins under unified reserves, and under unified reserves pays no more than they pay under personalized reserves.

43. **Proposition:** Consider a list of n bidders who face personalized reserves $r_1,\ldots,r_n$. Let also $r \geq r_i$ be a unified reserve. For any values $v_1,\ldots,v_n$, if Bidder i wins the item in a second-price auction with unified reserve r and pays p, then Bidder i also wins the item in a second-price auction with personalized reserves $r_1,\ldots,r_n$ and pays at least p.

44. **Proof:** The winner of a second-price auction with unified reserve r is the highest bidder whose value exceeds r, and the payment is their minimum bid to win. If this is Bidder i, it means that $v_i \geq r$, and also that v_i is the highest bidder. If no other bidder exceeds r, then Bidder i's minimum bid to win is r. If the highest other bidder exceeds r, then the highest other bid becomes Bidder i's minimum bid to win.

45. Let us now consider the same values in a second-price auction with personalized reserves $r_1,\ldots,r_n$. Because $r_i \leq r$, and Bidder i is the highest bidder, and $v_i \geq r$, this means that Bidder i still wins under personalized reserves, as desired. Moreover, the price paid by Bidder i is now equal to the maximum of r_i and the highest other bid that exceeds its reserve. In particular, if the highest other bid exceeds r, then under unified reserves Bidder i pays the highest other bid, and under personalized reserves Bidder i pays *at most* the highest other bid (because r_i is less than r is less than the highest other bid, and the highest other bid that clears its reserve is also less than the highest other bid. Bidder i's payment is the maximum of two terms that are both less than the highest other bid, so its payment is less than the highest other bid). If the highest other bid falls below r, then Bidder i pays r under unified reserves, and pays at most r under personalized reserves (because r_i is less than r, and the highest other bid that clears its reserve is less than r. Bidder i's payment is the maximum of two terms that are both less than r, so its payment is less than r).

46. Therefore, we conclude that whenever Bidder i would win under unified reserves, Bidder i surely wins under personalized reserves, and under personalized reserves pays no more than they would pay under unified reserves.

47. **Connecting second-price auctions to first-price auctions with efficient outcomes:** The two propositions above analyze the impact of second-price auctions with unified versus personalized reserves on bidders, but the body of the section instead concerns first-price auctions with efficient outcomes. Interestingly, these claims for second-price auctions immediately carry over to first-price auctions with efficient outcomes in two possible ways.

48. First, if bidders behave according to the Bayes correlated equilibrium proposed by Jin and Lu (2023), then they are essentially turning the first-price auction into a second-price auction (that is, bidders submit bids so that resulting outcome is identical as if all bidders submitted their true values to a second-price auction). Therefore, in the Bayes correlated equilibrium proposed by Jin and Lu (2023), the exact same outcomes arise as if bidders instead participated in a second-price auction.

49. Second, if bidders happen to bid in a Bayes-Nash equilibrium that is also efficient, then Myerson's revenue equivalence theorem establishes that the same bidder wins the item in this efficient Bayes-Nash equilibrium as in the second-price auction where each bidder bids truthfully, and in expectation the payments are the same too.

50. I have already noted theory alone is insufficient to confidently predict bidding behavior in first-price auctions, and justified earlier in this appendix that efficient outcomes are a natural assumption to understand the likely impact of reserve-price setting on bidders. The previous two paragraphs justify why analysis of second-price auctions extends to analysis of first-price auctions under efficient outcomes.

4) Impact of ad quality on reserve-setting heuristics

51. Here, I establish that increased ad quality results in lower reserves under both the treat-as-single heuristic and the eCPM heuristic. One should reasonably conclude that natural reserve-setting heuristics will set lower personalized reserves on bidders with higher ad quality.

52. **Proposition:** Consider two exchanges that have identical historical bids, but non-identical historical ad quality. That is, Exchange A has an average historical bid of b and an average historical ad harm of q_A (that is, q_A dollars of harm is caused on average), while Exchange B has an average historical bid of b and an average historical ad harm of $q_B>q_A$. Then the eCPM heuristic sets a higher personalized reserve on Exchange B than Exchange A (by exactly q_B-q_A).

53. **Proof:** The eCPM heuristic sets a personalized reserve of $b+q_A$ on Exchange A and $b+q_B$ on Exchange B. The personalized reserve set for Exchange B is therefore greater than that for Exchange A by exactly q_B-q_A.

54. **Proposition:** Consider two exchanges for which the publisher's prior on future bids is identically D, but Exchange A has a typical ad harm of q_A, while Exchange B has a typical ad harm of $q_B > q_A$. Then the treat-as-single heuristic sets a higher personalized reserve on Exchange B than Exchange A.

55. **Proof:** The treat-as-single heuristic sets a reserve r_A that optimizes $(r-q_A)$*Pr[bid from D exceeds r] for Exchange A, and r_B that optimizes $(r-q_B)$*Pr[bid from D exceeds r] for Exchange B. A well-known corollary of (Myerson 1981) establishes the manner by which one optimizes reserves for values drawn from D and a cost of c upon sale as the following: (a) compute the revenue curve, which plots $REV(q):=q*F_D^{-1}(q)$ as a function of q on domain [0,1], (b) compute IREV, the "upper concave envelope" of REV, which is concave, (c) find any point r for which the left-derivative at r is at least c and the right-derivative at r is at most c and for which IREV(q) = REV(q), then r is an optimal reserve (there may be multiple optima).

56. Because IREV is a concave function, the left- and right-derivatives are both increasing in q. Therefore, any q that satisfies the optimality condition for q_A is at least as large as any q' that satisfies the optimality condition for q_B (because the optimality condition for q_B requires a larger derivative than q_A). Therefore, any optimal reserve for ad harm q_A sells with higher probability than any optimal reserve for ad harm q_B, and therefore any optimal reserve itself for ad harm q_A must be lower than any optimal reserve for ad harm $q_B > q_A$.

APPENDIX G. CONDUCT ANALYSIS: DYNAMIC REVENUE SHARING

57. Here, I present opinions on the role of each component of the DRS conducts, as well as the isolated impacts of each of these components. These isolated impacts come together to constitute the overall impacts I provide in Section VII. The components I analyze are:

a. The impact of potential bid shading.

b. The impact of clearing the impression in the dynamic range.

c. The impact of clearing publisher and advertiser debt.

1) Impact of bid shading

58. Neither DRSv1 nor DRSv2 are truthful auctions. However, DRSv1 was not disclosed. Therefore, because AdX is described as a second-price auction with reserve (which is truthful), I would expect advertisers to bid their true values into DRSv1, and the impact of bid shading (if it occurred at all) to be minimal.

59. Aspects of DRSv2 were disclosed, and therefore advertisers might have responded by shading their bids. There are two ways advertisers might have responded: (a) an advertiser whose value lies in the dynamic region might have shaded their bid, but still submitted a bid, or (b) an advertiser whose value lies in the dynamic region may have forgone submitting a bid entirely. Perhaps surprisingly, bid-shading of form (a) has no impact on DRSv2 outcomes as compared to truthful bidding, due to DRSv2's debt collection mechanism and the fact that DRSv2 results in a debt-aware second-price mechanism to advertisers (meaning that all bids in the dynamic region are treated identically).[15] Advertiser behavior of form (b) results in identical outcomes as a clean second-price auction with no DRSv2 at all. That is, if all advertisers took behavior (b), which I previously noted is optimal, DRSv2 would result in identical outcomes to no DRS and the entire project would have been moot. To the extent that DRSv2 has any impact on anything, it is because some advertisers are bidding in the dynamic region. Therefore, the only potential impact that either (a) or (b) could have on DRSv2 is that it activates less often, but neither (a) nor (b) can possibly 'counter' any impact of DRSv2.

[15] The only potential impact is that bid-shading would pay less now and accumulate more debt, which might cause DRSv2 to activate less later.

60. In sum, the impact of bid shading on DRSv1 is likely minimal or non-existent, given that DRSv1 was not disclosed and AdX was described as a (truthful) second-price auction with reserve. Bid shading may have occurred in DRSv2, due to some aspects being disclosed, but such bid shading cannot possibly counteract any of the impacts discussed below.

2) Impact of clearing the impression in the dynamic range

61. When an impression is cleared in the dynamic region due to DRS, the publisher gains exactly its price floor, and loses exactly its opportunity cost had AdX not transacted the impression. Therefore, this acts as an indeterminate force on publishers' revenue under DRSv1/v2 as compared to no DRS. If the opportunity cost is on average lower than the reserve price, then publishers would see increased revenues under DRSv1/v2 as compared to no DRS. If the opportunity cost is on average higher than the reserve price, then publishers would see decreased revenue under DRSv1/v2 as compared to no DRS. It is worth acknowledging that a fully-informed revenue-maximizing publisher (both in terms of data on likely bids, and in terms of the complete auction format executed) should set a price floor exceeding their opportunity cost. Still, such publishers do not see increased revenue on a per-impression basis.

62. DRSv2 transacts in the dynamic region more often than DRSv1. Therefore, DRSv2 amplifies the impact of DRSv1 over no DRS (which is indeterminate without considering the opportunity cost of uncleared transactions, and I discuss above the manner by which to determine its impact). DRSv2 adds a mechanism to increase the average take-rate in comparison to DRSv1 on later transactions, which allows DRSv2 to transact in the dynamic region more often.

63. Bid-shading and transactions in the dynamic region are the only two impacts on publishers when comparing DRSv1 to no DRS. Assuming that advertisers did not shade their bids under DRSv1, because it was not disclosed, the impact on publishers is therefore completely determined by a comparison between their reserves and opportunity cost on impressions that clear through AdX in the dynamic region.

3) Impact of clearing debt

64. Under DRSv2, AdX keep per-publisher and per-advertiser debt accounts, which allowed for both a dynamic increase and decrease of the AdX take rate. This is in contrast to DRSv1, where there are no debt accounts, and the take rate is only can be lowered. Advertiser debt is cleared when AdX dynamically increases the take rate for an advertiser. If no advertiser debt were

cleared, then the advertisers would be indifferent between DRSv1 and DRSv2, since debt clearing is the only functional difference between DRSv1 and DRSv2.[16] However, all else being fixed, if any amount of advertiser debt is cleared, advertisers are worse off as a result.

65. Under DRSv2, if all advertiser debt is cleared, then publishers as a whole ultimately receive payments equal to their reserve prices on all impressions transacted in the dynamic region. If any advertiser debt remains uncleared, then publishers ultimately receive payments less than their reserve price on all impressions transacted in the dynamic region. If this were to happen, this could be viewed as publishers paying in order to implement DRSv1 more often.

66. Under DRSv2, clearing publisher debt (and isolating the impact of clearing debt) acts as a force towards decreasing publisher revenue under DRSv2 in comparison to DRSv1. This conclusion simply isolates the impact of publishers paying AdX additional money to clear debt and does not account for the impact of transacting additional impressions, which is covered above.

67. Under DRSv2, clearing advertiser debt (which has no impact on whether transactions are cleared in the dynamic region) acts as a force towards increasing publisher revenue under DRSv2 in comparison to DRSv1. Clearing advertiser debt is a pure transfer from advertisers to AdX, of which 80% goes to publishers. Therefore, clearing advertiser debt is a force towards increasing publisher revenue under DRSv2 in comparison to DRSv1.

68. Whenever a transaction is cleared in the dynamic region, the increase in publisher debt is exactly 80% of the increase in advertiser debt. Therefore, if all debts are cleared, publishers as a whole are ultimately paid 80% of advertiser payments to AdX. Therefore, if all debts are cleared across a billing period, the total payment to publishers as a whole from AdX across that billing period is exactly the same as if each transaction that cleared in the dynamic region paid the publishers their reserve price after a 20% take-rate from AdX. This further implies that the total payment to publishers from AdX across that billing period is exactly the same as if each

[16] This assumes that both DRSv1 and DRSv2 charge winners in the dynamic region their bid. As I noted, the documentation I have access to considered alternative possible payments for DRSv2 bidders in the dynamic region. In case these payment rules are used instead, all of my analysis here holds verbatim after replacing 'debt' with 'debt incurred even after charging the winner their bid'. That is, I consider the first step in DRSv2 to be 'charge the winner their bid eventually, whether now or via debt', and the second step to be 'further charge the winner remaining debt to hit the effective reserve'.

impression cleared according to a regular second-price with take rate of 20%.[17] If all publisher debt clears, publishers ultimately pay all publisher debt. If all advertiser debt clears, publishers ultimately receive 80% of all advertiser debt. Because publisher debt increases by exactly 80% of the increase in advertiser debt, this means that when all debt clears, publishers have indeed received 80% of all payments by Advertisers to AdX. As I noted above, when advertiser debt clears, advertisers ultimately pay the reserve price plus the AdX margin for any impressions they win in the dynamic region, and 80% of this is exactly the reserve price. The remaining payments outside the dynamic region are exactly 80% of what the advertiser pays in a regular second-price auction with a take rate of 20%, plus debt-clearing operations.

69. If all advertiser and publisher debt clears (and not yet accounting for which publisher receives 80% of cleared advertiser debt), the impact of debt tracking/clearing on publishers cancels out. Therefore, if all advertiser and publisher debt clears, the joint impact of DRSv2 over DRSv1 of debt tracking/clearing and transacting additional impressions amplifies the impact of DRSv1 over no DRS on publishers. If all publisher and advertiser debt clears, I noted above that the publisher payout is exactly equal to 80% of the transacted impressions, as if they were auctioned according to a regular second-price auction with take rate 20%. This means that the ultimate contribution to their transactions that clear in the dynamic region is exactly the reserve price they set. Therefore, assuming all advertiser and publisher debt clears, the ultimate comparison for publisher revenue when comparing DRSv2 to DRSv1 is the reserve prices on additional impressions cleared in the dynamic region versus the opportunity cost of these impressions clearing through AdX, which is the same as the comparison of DRSv1 to no DRS.[18]

70. Under DRSv2, if any advertiser debt remains uncleared, this acts as a force toward decreasing publishers' revenue when comparing DRSv2 to DRSv1. Similarly, under DRSv2, if any publisher debt remains uncleared, this acts as a force towards increasing publisher's revenue when comparing DRSv2 to DRSv1.

[17] Note that, of course, some of these impressions were cleared in the dynamic region. Therefore, the advertiser did not consider themselves to be participating in a regular second-price auction with a take rate of 20%, but the publisher ultimately views the outcome as such. The advertiser instead views it as a debt-aware second-price auction.

[18] Note that because the cleared impressions are different, it is possible that the average opportunity cost and reserve prices of impressions cleared in DRSv2 but not DRSv1 differ from those cleared in DRSv1. There is *a priori* no reason to expect the averages to lean one way or the other

71. Under DRSv2, an advertiser might incur a debt while purchasing an impression from one publisher and clear it while purchasing an impression from another publisher.[19] This aspect of DRSv2 therefore results in a transfer of funds among publishers but has no effect overall across publishers.

72. After isolating the impacts of various aspects of DRSv2, I can now opine on the role of each component and summarize the ways in which publishers, advertisers, and AdX are impacted.

a. Impact of clearing advertiser debt. Let me compare DRSv2 transacting exactly the same impressions but not tracking/clearing advertiser debt to DRSv2. In this case, tracking/clearing advertiser debt acts as a pure transfer from advertisers to AdX and publishers (20% goes to AdX, 80% goes to publishers). If no advertiser debt is cleared, then advertisers would be indifferent between DRSv2 and DRSv1.[20] If any advertiser debt is cleared, advertisers are harmed by DRSv2 in comparison to DRSv1. If all advertiser debt is cleared, then publishers ultimately receive payments equal to their reserves on all impressions transacted in the dynamic region. If any advertiser debt remains uncleared, then publishers ultimately receive payments less than their reserve prices across impressions transacted in the dynamic region. If this were to happen, I suggest viewing this as publishers paying in order to implement DRSv1 more often.

b. Impact of clearing publisher debt. When comparing DRSv2 transacting exactly the same impressions but not tracking/clearing publisher debt to DRSv2, tracking/clearing publisher debt acts as a pure transfer from publishers to AdX.

c. Impact of tracking/clearing debt at all. Finally, let me compare DRSv2 transacting exactly the same impressions but not tracking/clearing any debt at all to DRSv2. In this case, if exactly the same impressions were transacted, then each advertiser

[19] GOOG-NE-13207241 at -5. "AdX Dynamic Revshare v2: Launch Doc." (under the "Implementation of DRSv2" subsection)
The example on -6 demonstrates an advertiser b1 accumulating debt with one publisher p1 and clearing it with a different publisher p2.
[20] Again, this assumes that DRSv2 charged winners their bid in the dynamic region. As I previously described, if this is not the case, then my analysis considers the tracking of debt that would be owed even after advertisers eventually pay their bids (independent of whether this was done on the impression itself, or via debt).

would pay their bid on transactions cleared in the dynamic range, and each publisher would receive their reserve prices on such transactions. I have already noted above that when all debt clears that each advertiser ultimately pays the reserve price plus the AdX margin on all transactions that clear in the dynamic region, and publishers ultimately receive their reserve prices on such transactions. Therefore, the impact of tracking/clearing debt at all is simply a pure transfer from advertisers to AdX (assuming that all debts clear). In particular, AdX could run DRSv2 exactly as-is without tracking debt, and this would strictly help advertisers, not affect publishers at all, and generate less revenue for AdX (of course, AdX would still profit from DRSv2 without tracking debt over no DRS, as they still clear additional transactions).[21]

4) Comparison of DRS versions

73. The list below shows a chain of comparisons of no DRS, DRSv1, DRSv2, and hypothetical counterfactuals, where each step in the chain has a clear comparison. Below, I will use the phrase "DR-friendly" to denote publishers whose reserve exceeds their opportunity cost when transacting in the dynamic region, and "DR-unfriendly" to denote publishers whose opportunity cost exceeds their reserve when transacting in the dynamic region. I will also use the phrase "v2-lucky" to denote publishers who wind up clearing more advertiser debt than building it, and "v2-unlucky" to denote publishers who wind up building more advertiser debt than clearing it.

a. **No DRS versus DRSv1, assuming the advertisers bid truthfully**. In such a scenario, AdX sees increased revenues due to transactions cleared in the dynamic region. Advertisers are affected in terms of outcomes, but neutral in terms of payoffs.[22] DR-friendly publishers see increased revenues, DR-unfriendly publishers see decreased revenues, and DR-neutral publishers have no impact on revenues.

[21] It would ultimately help publishers, because advertisers would shade their bids less aggressively if they paid their bid on impressions won in the dynamic region as opposed to the reserve price plus the AdX margin, which exceeds their bid.
[22] If their values accurately capture the impact of budget constraints, they are truly neutral. If they further care about ROI, they may consider this a form of harm, because their return on investment decreases without seeing any impact in payoff.

b. **DRSv1 versus no-debt DRSv2, assuming advertisers bid truthfully**[23] AdX sees further increased revenues, because more transactions are cleared in the dynamic region. Advertisers are again affected in terms of outcomes but neutral in terms of payoffs.[24] DR-friendly publishers see further increased revenues and DR-unfriendly publishers see further decreased. DR-neutral publishers again have no impact on revenues.

c. **No-debt DRSv2 versus DRSv2 with per-publisher debt pools, assuming advertisers bid truthfully.**[25] AdX sees further increased revenues, because they receive 20% of all cleared advertiser debt. Advertisers lose payoff exactly equal to the cleared debt. All publishers are unaffected.[26] This is a pure transfer from advertisers to AdX; AdX benefits by exactly as much as advertisers are harmed.

d. **DRSv2 with per-publisher debt pools versus DRSv2, assuming advertisers bid truthfully**. AdX and advertisers are both unaffected (as in, they see exactly the same outcomes.[27] v2-lucky publishers see increased revenues, v2-unlucky publishers see decreased revenues, and v2-neutral publishers are unaffected This is a pure transfer from v2-unlucky publishers to v2-lucky publishers, the entire process is zero-sum.

e. **DRSv2, assuming advertisers bid truthfully versus DRSv2, advertisers shade their bids**. If advertisers shade their bids in the dynamic region, the outcomes compared to DRSv2 with truthful bidding are unchanged.[28] If an advertiser sometimes foregoes bidding in the dynamic region entirely, DRSv2 reverts to no DRS. Therefore, the only potential impact of bid-shading is to lessen each of the impacts in the previous steps but bid-shading cannot possibly reverse any step. I

[23] No-debt DRSv2 refers to the hypothetical scenario where AdX transacts exactly the same impressions as DRSv2, but just ignores the debt-clearing transactions.
[24] And again might consider themselves to be harmed because their ROI goes down without seeing an increased payoff.
[25] This means that exactly the same transactions are cleared as DRSv2, except that when advertiser debt is cleared, 80% of it is paid to the publisher with whom it was generated, instead of the publisher from whom the debt-clearing impression is purchased.
[26] Assuming all debt clears.
[27] The decision to do a single debt pool for each advertiser, rather than each (advertiser, publisher) pair allows more debt to clear, and so would enable more applications of DRSv2. This would cause increased revenues for AdX.
[28] As previously noted, the only possible change is that the advertiser incurs more debt that may take time to clear, and so DRSv2 may activate less often. This can lessen the impact in the previous steps, but cannot reverse any step.

again note that if all advertisers forewent bidding in the dynamic region entirely, then DRSv2 would be completely moot and execute exactly as no DRS.[29]

74. The above chain helps to process the entire impact of DRSv2 on various parties.

a. Only a DR-friendly/neutral and v2-lucky/neutral publisher has non-negative revenue impact in all transitions. Any publisher that is either DR-unfriendly or v2-unlucky would need to quantitatively compare the impacts at different transitions to determine their impact from DRSv2.

b. AdX strictly benefits in each of the first three transitions. In particular, the decision to track and clear debt results in a pure transfer from advertisers to AdX, on top of an already-profitable program for AdX.

c. Advertisers cannot possibly benefit under DRSv2 as compared to no DRS. Any advertiser that ever wins an auction when their value lies in the dynamic region suffers a loss in payoff compared to no DRS by ultimately paying more than their value for the impression won. Moreover, if no advertiser wins an auction when their value lies in the dynamic region, then DRSv2 is completely moot and equivalent to no DRS.[30] Advertisers who continued to bid truthfully under DRSv1 were neutral in terms of payoff, and negative in terms of ROI.[31]

5) Comparison of outcomes under DRS versions

75. In Tables 2 through 6 below, I provide a summary of various outcomes of the auctions under no DRS, DRSv1 and DRSv2. Each row corresponds to a specific instance of an auction outcome depending on the relationship between the highest and the second highest bids, as well as the reserve price.

[29] This suggests that if all advertisers took this action, the entire DRSv2 program would have been mot and serve no purpose. Moreover, if the program were entirely moot, it could not possibly have generated increased revenue for any publishers, whereas GOOG-NE-13207241 claims a revenue lift. GOOG-NE-13207241. “AdX Dynamic Revshare v2: Launch Doc.”

[30] Again, I note that if DRSv2 led to an increase in AdX revenue, that would strongly suggest that a significant volume of advertisers must have bid in the dynamic region, against their self-interest.

[31] An advertiser who became aware of DRSv1, despite the fact that it was concealed, may have been able to profitably shade their bids. If so, such an advertiser could have benefit under DRSv1 compared to no DRS.

76. The tables utilize the following notation: r denotes the reserve price set to AdX. b_1 denotes the highest bid submitted to AdX and b_2 denotes the second highest bid submitted to AdX. x and y denote non-negative numbers that vary to clear pools. c denotes the opportunity cost if impression is not sold to AdX. c could be 0, if impression otherwise unsold, or r, if r comes from header bidding, or anything higher than 0, if it denotes the expected revenue achieved throughout the rest of the waterfall.

Table 1: Summary of various outcomes of the auctions under DRSv2.

Mathematical Condition	DRSv2 Revenue (amount paid by advertiser to AdX)	DRSv2 Payout (amount paid by AdX to publisher)	Relative Margin ((Revenue – Payout)/Revenue)	Change in publisher's debt	Change in advertiser's debt
$b_1<r$, highest bidder does not clear price floor.	0, because impression unsold.	0, because impression unsold.	20%	0	0
$b_2 \geq 1.25r$, second-highest bid clears price floor with the AdX margin	b_2+x, second-highest bid plus the amount of debt to clear.	$0.8\ b_2 + 0.8\ x – y$.	20% when $y = 0$, more than 20% when $y > 0$.	$-y$, because publisher pays off debt.	$-x$, pays x beyond clearing price.
$1.25r > b_1 \geq r$	b_1, because this is the dynamic region.	r, because this is the dynamic region.	$(b_1-r)/b_1 < 20\%$, because $b_1 < 1.25r$.	$+r - 0.8\ b_1$, because payout is higher than 80% of the revenue, so publisher incurs debt.	$+1.25r-b_1$, because buyer got impression for b_1 that "should have cost 1.25r".
$b_1 \geq 1.25\ r > b_2$, highest bid clears price floor with the AdX margin	$1.25r + x$, price floor with the AdX margin plus debt payment.	$r + 0.8x - y$	20% when $y = 0$, greater than 20% when $y > 0$.	$+y$, because publisher pays off debt.	$-x$, pays x beyond clearing price.

Table 2: summary of various outcomes of the auctions under DRSv1.

Mathematical Condition	DRSv1 Revenue (amount paid by Advertiser to AdX)	DRSv1 Payout (amount paid by AdX to Publisher)	Relative Margin ((Rev – Payout)/Revenue)	Change to Publisher's absolute margin from 20%.
$b_1 < r$, highest bidder does not clear price floor.	0, because impression unsold.	0, because impression unsold	20%	0
$b_2 \geq 1.25p$, second-highest bid clears price floor + margin	b_2, second-highest bid.	0.8 b_2	20%	0
$1.25r > b_1 \geq r$	b_1, because this is the dynamic region.	r, because this is the dynamic region.	$(b_1-r)/ b_1 < 20\%$, because $b_1 < 1.25r$.	0.8 b_1 - r, because payout is greater than 80% of revenue, so publisher absolute margin goes further down from 20%.
$b_1 \geq 1.25 r > b_2$, highest bid clears price floor with the AdX margin	1.25r, price floor with the AdX margin	r	20%	0

Table 3: Summary of various outcomes of the auctions under no DRS, DRSv1 and DRSv2.

Mathematical Condition	No DRS Revenue (amount paid by advertiser to AdX)	No DRS Payout (amount paid by AdX to publisher)	No DRS Income (received by publisher from any source)	DRSv1 Income (received by publisher from any source).	DRSv2 Income (received by publisher from any source).
$b_1 < r$, highest bidder does not clear price floor.	0, because impression unsold.	0, because impression unsold	c, sold elsewhere.	c, sold elsewhere.	c, sold elsewhere.
$b_2 \geq 1.25p$, second-highest bid clears price floor + margin	b_2, second-highest bid.	0.8 b_2	b_2	b_2	$b_2+.8x-y$
$1.25r > b_1 \geq r$	0, not sold to AdX	0	c, sold elsewhere	r	r
$b_1 \geq 1.25\ r > b_2$, highest bid clears price floor with the AdX margin	1.25r, price floor with the AdX margin	r	r	r	r+.8x-y

Table 4: Summary of various outcomes of the auctions under no DRS, DRSv1 and DRSv2.

Mathematical Condition	No DRS Income (received by publisher from any source)	DRSv1 Income (received by publisher from any source).	DRSv2 Income (received by publisher from any source).	Change in Publisher's DRSv2 Debt	Change in Advertiser's DRSv2 Debt
$b_1 < r$, highest bidder does not clear price floor.	c, sold elsewhere	c, sold elsewhere	c, sold elsewhere	0	0
$b_2 \geq 1.25p$, second-highest bid clears price floor + margin	b_2	b_2	$b_2+0.8x-y$	-y, because publisher pays off y debt.	-x, pays x beyond clearing price.
$1.25r > b_1 \geq r$	c, sold elsewhere	r	r	$+r-0.8\ b_1$, because payout is greater than 80% of revenue, so publisher incurs debt.	$+1.25r-b_1$, because buyer got impression for b_1 that "should have cost 1.25r."
$b_1 \geq 1.25\ r > b_2$, highest bid clears price floor with the AdX margin	r	r	r+.8x-y	-y, because publisher pays y off y debt.	-x, pays x beyond clearing price.

Table 5: Summary of various outcomes of the auctions under no DRS, DRSv1 and DRSv2.

Mathematical Condition	No DRS AdX Advertiser Payoff	DRSv1 AdX Advertiser Payoff	DRSv2 AdX Advertiser Payoff	Change in Publisher's DRSv2 Debt	Change in Advertiser's DRSv2 Debt
$b_1 < r$, highest bidder does not clear price floor.	0	0	0	0	0
$b_2 \geq 1.25p$, second-highest bid clears price floor + margin	$b_1 - b_2$	$b_1 - b_2$	$b_1 - b_2 - x$	-y, because publisher pays off y debt.	-x, pays x beyond clearing price.
$1.25r > b_1 \geq r$	0	0	0	$0.8\ b_1 - r$, because payout is greater than 80% of revenue, so publisher takes from the debt pool.	$+1.25r - b_1$, because buyer got impression for b_1 that "should have cost 1.25r."
$b_1 \geq 1.25\ r > b_2$, highest bid clears price floor with the AdX margin	b_1-1.25r	b_1-1.25r	b_1-1.25r-x	-y, because publisher pays off y debt.	-x, pays x beyond clearing price.

APPENDIX H. CONDUCT ANALYSIS: PROJECT BERNANKE

1) Collusion in Second-Price Auctions

77. Imagine a large group of friends are all interested in purchasing a first edition Charizard Pokémon card in a second-price auction with reserve on eBay. If each person is interested in optimizing their own payoff and only their own payoff, they should each bid their true value in the auction, because second-price auctions with reserve are truthful. But imagine instead that they care about each other's payoffs as well, since they are friends, and as a result, try to increase their collective payoff.

78. Auction theory shows that they can increase their collective payoff. First, they figure out who has the highest value for the item. Say it is Bidder One, and their value is *v*. Then Bidder One should submit a bid of *v* to the auction, and all other bidders should submit a bid of 0. By lowering their bids, they do not lower the chance of someone from your group winning the item, but they might lower the price Bidder One pays if they win.

79. For example, imagine that Bidder One's value is $300, their friends have values of $280, $260, $250, and $200, and the maximum value outside of the friend group is $250. If everyone bids truthfully (which maximizes their individual payoffs), the submitted bids will be $300, $280, $260, $250, $200, and $250, so Bidder One will win the item and pay $280. If instead the friends collude, the submitted bids will be $300, $0, $0, $0, $0, $250, so Bidder One will win the item and pay $250. This causes the friend group to gain $30 and causes the seller to lose $30. In fact, any group that wishes to collude in a second-price auction can optimally collude by submitting a single bid equal to the maximum value in that group and decreasing others' bids to $0. Hence, optimal collusion is straightforward among trusted parties.

80. Now imagine that the parties are not close friends but classmates who attend your same high school. They are happy to help each other over a random seller on the internet, but you do not necessarily trust each other at the same level and do not know each other's values. Imagine the same valuation as the previous example and let Bidder Two be the one with the value of $260. Since the maximum outside bid is $250, which is below their value, Bidder Two would like to win. So, Bidder Two would lie and tell their classmates that their value is in fact $500. They would optimally collude, and the submitted bids will be $0, $0, $500, $0, $0, $250, which causes Bidder

Two to win the item for \$250 and gain \$260−\$250=\$10.[32] This example highlights that collusion is tricky among untrusted parties.

81. To more effectively collude among untrusted parties, one solution is to turn the collusion itself into a second-price auction. Specifically, it can be proposed that any benefit the colluder gains by the collusion must be diverted to a third party. It could be implemented as follows:

a. Let the bidders submit bids to a group-wide auction, and let b^c_2 denote the second highest bid among the colluders,

b. Let the highest bidder submit a bid to the real second-price auction, and the rest of the colluders will bid zero,

c. When the real second-price concludes, if the colluder wins at a price higher than b^c_2, the colluders observe that the collusion did not benefit them in the end, and so the process concludes.

d. If the colluder loses, then the collusion again did not help them, so the process concludes.

e. If instead, the colluder wins at a price p lower than b^c_2, the group recognizes that the colluder saved b^c_2-p, and the winning colluder pays this amount to the collusion pool.

82. Through such a mechanism, each of the colluders now participates in a true second-price auction, even at the colluding stage. A colluder wins if and only if they are the highest bidder among all bidders, and if they win, they pay the second-highest bid among all bidders. The only difference is that the colluders pocket some revenue that would have otherwise gone to the seller, since the seller only gets the highest bid outside of the collusion, and the colluders keep any revenue beyond that. This shows that it is possible to run a truthful collusion process that allows the colluders to siphon off revenue from the seller to a third-party. But in order for this process to

[32] If my classmates do not audit Bidder Two's behavior in the auction, Bidder Two would actually do even better by claiming that they have a value of \$500, just to get the other classmates to bid \$0, but then still bidding Bidder Two's true value of \$260 in the second-price auction to be safe.

be truthful, each colluder must not themselves directly derive utility when this third-party receives money.[33]

83. In sum, I demonstrate that (a) collusion among bidders is profitable, (b) collusion is easy among bidders who fully trust each other and (c) collusion is hard among bidders who are strategic amongst themselves, but possible if they give the colluding profits to a third-party instead of themselves.

2) Overbidding in Second-Price Auctions

84. As demonstrated in Section II, bidding one's true value in a second-price auction is optimal. This means that overbidding (bidding above one's true value) is not optimal. To see why, imagine that a bidder's true value is $10, and they are thinking of submitting a bid of either $10 or $15. If the highest other bid *b* is less than $10, then the bidder will win with either bid and pay *b*. If the highest other bid *b* is more than $15, then the bidder will lose with either bid. If instead the highest bid *b* lies between $10 and $15, the bidder will lose with a bid of $10 but win with a bid of $15 and pay *b* which is higher than $10. That is, the only change a bidder can possibly cause by overbidding is to win an item when they are not the highest value bidder and pay more than their value for the item. This would lead to a net loss for this bidder.

85. Overbidding also has an impact on other parties in an auction. Let v denote the overbidder's value, v_1 denote the highest value (which is higher than v), and v_2 denote the second-highest value (which is either higher than v, or equal to v if the overbidder is the second highest bidder). If the overbidder does not overbid more than v_2, then it is as if they never overbid, and no change occurs. If the overbidder overbids more than v_1, then (a) the overbidder now wins and pays v_1 and suffers a loss of $v_1 - v$ for paying v_1 for an item they value at v, (b) the seller enjoys revenue v_1 instead of v_2, and (c) the highest-value bidder suffers a loss of $v_1 - v_2$, they previously won the item and paid v_2, but now get nothing. This means that all parties together see decreased payoff as a group. The payoff decrease of the previous winner cancels out the increased payoff to the seller, and the payoff decrease of the overbidder is strictly positive. Furthermore, the seller and overbidder together do not get increased payoff (but may have the same joint payoff as

[33] Note that this is not the mechanism that Project Bernanke or Project Global Bernanke uses. Instead of charging the winning colluder their minimum bid to win, Project Bernanke and Global Bernanke charges their bid. This successfully disincentivizes the overbidding seen in the example but might incentivize bid-shading. GOOG-DOJ-28385887 proposes exactly this truthful procedure as a future direction, but I do not know whether it was later implemented. GOOG-DOJ-28385887 at -03. August 17, 2015. "Beyond Bernanke."

before). The seller's gain of $v_1 - v_2$ matches the overbidder's loss of $v_1 - v$ if and only if $v = v_2$. If v is less than v_2, then the overbidder's loss strictly exceeds the seller's gain.

86. The remaining possibility is that the overbidder overbids to a bid b more than v_2 but less than v_1. In this case: (a) the overbidder suffers no decrease in payoff, as they still lose, (b) the seller enjoys revenue of b instead of v_2, for a gain of $b-v_2$, and (c) the highest-value bidder now pays b instead of v_2, and therefore suffers a payoff decrease of $b-v_2$. I would like to highlight that all parties *together* are neutral – the payoff decrease of the highest-value bidder is exactly cancelled out by the increased revenue to the seller. I would also like to highlight that the seller and overbidder together benefit – the seller gets extra revenue due to the overbid, and the overbidder suffers no decrease in payoff.

87. Let me also explain what might happen in case the item was previously unsold due to a high reserve. In this case, again let v denote the overbidder's value, and r denote the reserve (which is > v). If the overbidder does not overbid more than r, then it is as if they never overbid, and no change occurs. If the overbidder overbids more than r, then: (a) the overbidder now wins and pays r and suffers a payoff decrease of $r - v$ for paying r for an item they value at v, and (b) the seller enjoys revenue of r instead of 0. In this case, I would like to highlight that all parties together see a joint payoff increase of v (intuitively, this is because the payments made between overbidder and seller are zero-sum, and now the overbidder gets an item at value v whereas previously the item was unsold).

88. In sum, when there is overbidding:

a. The overbidder never increases their payoff and sometimes decreases their payoff (if they overbid enough to win).

b. The previous highest bidder experiences decreased payoff (as long as the overbid is high enough to have any impact on the outcome at all).

c. The seller sees increased revenues (as long as the overbid is high enough to have impact on the outcome at all).

d. The entire ecosystem sometimes sees increased joint payoff (if the item was previously unsold) and sometimes sees decreased joint payoff (if the item was previously sold, and the overbidder bids high enough to win).

e. The overbidder and seller jointly might see increased payoff (if the overbidder bids enough to raise the second highest bid but not enough to win, or if the item was previously unsold and is now sold to the overbidder), might see decreased payoff (if the overbidder was previously not the second-highest bid, and still overbid enough to win), and might have unchanged payoff (if the overbidder was previously the second-highest bid, and overbid enough to win, or if the overbid is too low to affect the outcome).

f. The previous highest-bidder and seller jointly have the same payoff, because the seller sees increased payoff by exactly the payoff decreased of the previous highest-bidder.[34]

3) Collusion in First-Price Auctions

89. Imagine again that a large group of friends are all interested in purchasing a first-edition Charizard Pokémon card, but this time the auctioneer runs a first-price auction with reserve. If each person is interested in optimizing their own payoff, they should each form a probabilistic belief regarding the highest other bid and submit the bid that maximizes their profit (their value minus their bid) times the probability their bid exceeds the highest other bid.

90. But imagine instead that they care about their friends. It turns out that they can jointly maximize their payoffs by figuring out who has the highest value for the item. Then they form a probabilistic belief regarding the highest bid *outside of the friend group*, and friends who do not have the highest value submit a bid of $0. By lowering their bids, they reduce the uncertainty regarding the highest other bid, and this allows the highest bidder to shade their bids more aggressively.

[34] To further see this last point, think of the overbidder's bid as continuously ticking up from v. Until it ticks up to v_2, there is no impact. As it ticks up from v_2 to v_1, it continuously drains money from the highest bidder to the auctioneer, because it continuously increases the price the highest bidder must pay. Once it ticks to v_1, the overbidder now wins, the highest bidder sees a payoff decrease of v_1-v_2 (because now they get nothing, whereas previously they got the item and paid v_2) and the auctioneer sees increased revenue of v_1-v_2 (because now they are paid v_1, whereas previously they got paid v_2).

91. To see this with an example, imagine that a Pokémon enthusiast group of friends form six of the ten bidders in the auction. Moreover, they believe that everyone's value is equally likely to be any number between 0 and 500.[35] Then the equilibrium[36] of a first-price auction where no bidders collude is for each bidder to shade their bid by exactly 10%. That is, if a bidder has value *v*, they should shade their bid to 9*v*/10 and expect all other bidders to share their bid by the same 10%.[37] So without collusion, each of your friends would shade their bid by 10%. With collusion, the friends could get together and first determine who has the highest value among the group. Whoever has the highest value is now only competing against four other bidders, rather than nine. In this case, assuming that the other bidders all shade their bid by the same margin (which might be 10%, if they do not know the friends are colluding), the winner of your group's optimal bid is to shade more aggressively, by 20%. This example highlights that collusion is possible in a first-price auction, optimal collusion results in more aggressive bid shading, and this aggressive bid shading is enabled by sharing knowledge among colluders to remove the need to bid high enough to outbid each other.[38]

92. Like with a second-price auction, collusion among untrusted parties is more challenging. For example, perhaps the top two values among friends are $400 and $380. The optimal collusion is for the higher value friend to shade their bid from 400 down to 0.8*$400 = $320. But knowing this, the lower value friend might now just enter the auction anyway with a bid of $320.01. Still, this can be addressed with the optimizing devices I explained in the body of this section and a third party.

93. To see this with an example, imagine that six friends are at the same high school, and are aware of a math prodigy who herself has no interest in Pokémon cards, and they prefer that the math prodigy collect extra revenues from the auction than the anonymous seller, if possible.

[35] That is, the Bayesian prior everyone forms over others' values is that each value is drawn independently from the uniform distribution on [0,500], as in the independent private values model.

[36] Here and throughout the report, I use the equilibrium notion called the "Bayes-Nash Equilibrium." Bayes-Nash equilibrium builds on the regular Nash equilibrium by incorporating incomplete information, and it is the standard equilibrium concept to analyze these types of auctions in the literature.

[37] In fact, with n bidders whose values are drawn uniformly on some interval [0,k], it is a Bayes-Nash equilibrium for each bidder to shade their bid by 1/n. That is, a bidder with a value of v should submit a bid of (1 − 1/n)*v and expect other bidders to shade their bids by the same ratio.

[38] As a more extreme example, imagine that your friends are the only bidders and there is no reserve. Then you should find the highest value friend and have them submit a bid of .01 to win the item. Collusion in a first-price auction is much harder to optimize than in a second-price auction, but in extreme cases can be just as powerful.

Without jointly colluding, they could ask the math prodigy to play the role of the bid optimizer. Their bidding function would be $b(v) := 9v/10$ for each of them. If any of the friends won, and the second highest submitted bid was w, the math prodigy would pay $9v/10$ to the auction and charge the winner $10w/9$ (the minimum bid they could have submitted and still submitted a bid exceeding w after shading by $b(\cdot)$). If instead the friends jointly colluded with the math prodigy, they could find the highest value v among them and submit a bid of $4v/5$ (the optimal bid-shading with only four other bidders is by 20%). If they win, let w denote the highest other bid in the auction and v_2 denote the highest other bid among the friends. The math prodigy would pay $4v/5$ to the auction and charge the winner the maximum of $5w/4$ and v_2. This is the minimum bid to win for the winner of the entire joint process. The joint winner needs to submit a bid such that it both wins the internal auction (and therefore exceeds v_2) and generates a winning bid in the external auction (and therefore exceeds $5w/4$). The key takeaway here is that joint collusion in first-price auction is possible, and with a third party who enforces the procedure.

4) Overbidding in First-Price Auctions

94. While first-price auctions are not truthful, overbidding in a first-price auction is still *always* suboptimal (bidding the true value, or anything less, is a better strategy). To see this, observe that overbidding guarantees a bidder non-positive utility. Either they lose and get nothing or win and pay more than their value. On the other hand, bidding their value (or anything less) guarantees a non-negative payoff. Either they lose and get nothing or win and pay less than their value. Like with a second-price auction, the only change a bidder can possibly cause by overbidding is to win an item and pay more than their value. Overbidding benefits the seller if it happens. If the previous highest bid is b, and a bidder overbids to b', the seller's revenue increases from b to b'. If the bidder's value is $v<b'$, they lose $b' - v$.

95. Consider overbidding in a first-price auction that was modified to be truthful as I discussed above. Recall that there are now three components in the auction: (a) the bidders, (b) the auction theorist who optimizes a bid on their behalf, and (c) the ultimate first-price auction. Through this lens, "truthful bidding" refers to the bidders reporting their true values to the auction-theorist, and the auction-theorist optimizing the bids for the first-price auction based on those values. Overbidding now refers to the auction-theorist optimizing a bid for the first-price auction based on a higher value.

96. Since in a first-price auction the bids are not going to be equal to the bidders' values (at least for the ones who are not colluding), it is much more involved to make quantitative claims about the joint impact of overbidding on multiple parties. But, it is certainly the case that overbidding (a) causes the overbidder a decrease in payoff (the only impact it can possibly have is to win an item and pay more than the value), (b) increases the seller's revenue (higher bids always lead to more revenue for the seller), and (c) causes a decrease in payoff for all other bidders (facing higher bids always leads to a lower win rate, and therefore lower utility).

5) Numerical Examples for Project Bernanke

97. Tables 1, 2 and 3 below work through several cases of possible bids, and computes the payments made by GDN advertisers to GDN, the payment made from GDN to AdX, and the counterfactual of what would have happened without Project Bernanke. Below, I summarize the key takeaways via examples.

a. **Example 1**. First, I provide an example where GDN still "second-prices itself" with Project Bernanke. Imagine that GDN has selected $\alpha = 4$ and $\beta = 0.25$, and GDN's highest bidder is $20 and its second highest bidder is $16. Then GDN submits bids of 4*$20=$80 and 0.25*$16=$4 to AdX. Imagine also that AdX's reserve is $2, and its highest other received bid is $3. Then GDN will win the auction, paying $4 to AdX because GDN's second highest bid becomes AdX's second highest bid. To compute the price charged to the highest bidder, GDN computes its second highest bid ($16), its highest bid ($20), and GDN's minimum bid to win on AdX ($3), and computes a price of max{$16,min{$20,$3}}=$16. That is, GDN charges its highest bidder $16, and pays $3 to AdX, for a margin of 13/16 = 81.25%. A 14% margin on $16 is $2.24, so GDN adds the difference of $13 − $2.24 = $10.76 to its Bernanke Pool.[39] If instead GDN did not use Project Bernanke, GDN would have submitted bids of $20*0.86 = $17.20 and $16*0.86 = $13.76 to AdX. GDN would have still won AdX's second-price auction with reserve $2 and paid $13.76 to AdX. It would have charged its highest bidder exactly its minimum bid to win, accounting for GDN's 14% take-rate, which is $16. In this case, we see that (a) all advertisers have the same payoff under Project Bernanke as with no Bernanke, (b) the

[39] Whether the pool is per-publisher or AdX-wide is immaterial to this section, so these examples work under both Project Bernanke and Project Global Bernanke.

publisher sees decreased revenue under Project Bernanke in this case since their revenue decreases from $13.76 without Project Bernanke to $3 with Project Bernanke, for a loss of $10.76. This is exactly how much is added to the Bernanke pool. This is a case where Project Bernanke is akin to collusion among GDN bidders, except the collusion is not optimal.

b. **Example 2**. Next, I give an example where GDN drops its second highest bid entirely from the AdX auction. Imagine that GDN has selected $\alpha = 4$ and $\beta = 0.25$, that GDN's highest bidder is $20 and its second highest bidder is $16. Then GDN submits bids of 4*$20 = $80 and 0.25*$16 = $4 to AdX. Imagine also that AdX's reserve is $2, and its highest other received bid is $10. Then GDN will win the auction, paying $10 to AdX because AdX's highest other bid becomes its second highest bid. To compute the price charged to the highest bidder, GDN computes its second highest bid ($16), its highest bid ($20), and GDN's minimum bid to win on AdX ($10), and computes a price of max{$16,min{$20,$10}} = $16. That is, GDN charges its highest bidder $16, and pays $10 to AdX, for a margin of 6/16 = 37.5%. A 14% margin on $16 is $2.24, so GDN adds the difference of $6 − $2.24 = $3.76 to its Bernanke Pool. If instead GDN did not use Project Bernanke, GDN would have submitted bids of $20*0.86 = $17.20 and $16*0.86 = $13.76 to AdX. GDN would have still won AdX's second-price auction with reserve $2 and paid $13.76 to AdX. It would have charged its highest bidder exactly its minimum bid to win, accounting for GDN's 14% take-rate, which is $16. In this case (a) all advertisers are have the same payoff under Project Bernanke as with no Bernanke, (b) the publisher sees decreased revenue under Project Bernanke because their revenue decreases from $13.76 to $10 with Project Bernanke, for a loss of $3.76. This is exactly how much is added to the Bernanke pool. This is a case where Project Bernanke is akin to collusion among GDN bidders, and the collusion is functionally optimal (because GDN lowers its second highest bid sufficiently below AdX's highest other bid).

c. **Example 3**. Next, I give an example where GDN wins an auction at a price above its highest bidder's value. Imagine that GDN has selected $\alpha = 4$ and $\beta = 0.25$, that GDN's highest bidder is $20 and its second highest bidder is $16. Then GDN

submits bids of 4*$20 = $80 and 0.25*$16 = $4 to AdX. Imagine that AdX's reserve is $30, and its highest other received bid is $40 (and that no other bidder clears the reserve). Then GDN will win the auction, paying $40 to AdX because AdX's highest other bid becomes its second highest bid. To compute the price charged to the highest bidder, GDN computes its second highest bid ($16), its highest bid ($20), and GDN's minimum bid to win on AdX ($40), and computes a price of max{$16,min{$20,$40}} = $20. That is, GDN charges its highest bidder $20, and pays $40 to AdX, for a margin of −20/20 = −100%. A 14% margin on 20 is $2.80, so GDN subtracts the difference of $2.80 + $20 = $22.80 from its Bernanke Pool. If instead GDN did not use Project Bernanke, GDN would have submitted bids of 20*0.86 = $17.20 and $16*0.86 = $13.76 to AdX. AdX's highest other bidder would have won, and paid $30. In this case (a) all GDN advertisers have the same payoff under Project Bernanke as no Bernanke (the now-winning GDN bidder went from losing to winning-and-paying-its-value, which give a payoff of 0), (b) the publisher sees increased revenue under Project Bernanke in this case since their revenue increases from $30 to $40 with Project Bernanke, for a gain of $10, which is less than is consumed from the Bernanke pool, (c) whatever non-GDN advertiser submitted a bid of $40 to AdX is sees a payoff decrease of $10, because they previously won the impression and paid $30, for a payoff of $10, but under Project Bernanke they instead lose for a payoff of 0. This is a case where Project Bernanke is akin to overbidding, but the overbidding decreases the bidder's payoff more than it increases the seller's revenue because GDN's highest value is lower than the price paid to the publisher without Project Bernanke.

d. **Example 4**. Finally, I give another example where GDN wins an auction at a price above its highest bidder's value. Imagine that GDN has selected α = 4 and β = 0.25, that GDN's highest bidder is $20 and its second-highest bidder is $16. Then GDN submits bids of 4*$20 = $80 and 0.25*$16 = $4 to AdX. Imagine that AdX's reserve is $30 but has not received any bids above the reserve. Then GDN will win the auction, paying $30 to AdX because AdX's reserve becomes its second highest bid. To compute the price charged to the highest bidder, GDN computes its second highest bid ($16), its highest bid ($20), and GDN's minimum bid to win on AdX ($30), and computes a price of max{$16,min{$20,$30}} = $20. That is,

GDN charges its highest bidder \$20, and pays \$30 to AdX, for a margin of −10/20 = −50%. A 14% margin on 20 is \$2.80, so GDN subtracts the difference of \$2.80 + \$10 = \$12.80 to its Bernanke Pool. If instead GDN did not use Project Bernanke, GDN would have submitted bids of \$20*0.86 = \$17.20 and \$16*0.86 = \$13.76 to AdX. AdX's auction would have failed to clear because no bids clear the reserve. In this case, we see that: (a) all advertisers have the same payoff under Project Bernanke as no Bernanke (the now-winning GDN bidder went from losing to winning-and-paying-its-value, which give an equal payoff of 0), (b) the publisher sees increased revenue under Project Bernanke in this case because their revenue increases from \$0 to \$30 with Project Bernanke, for a gain of \$30, which is more than is consumed from the Bernanke pool. This is a case where Project Bernanke is akin to overbidding, and the overbidding increases the seller's revenue more than it decreases the bidder's payoff because GDN's value is higher than the price paid to the publisher without Project Bernanke).

98. In Tables 7, 8 and 9 below, I utilize the following notation: c is the maximum of highest bid from AdX buyer and reserve price (*i.e.,* minimum bid to win for GDN). b_1 highest bid submitted to GDN and b_2 is the second highest bid submitted to GDN. α and β are Bernanke multipliers as explained in the report.[40]

[40] These tables are drawn from GOOG-AT-MDL-008881638, although I have lightly edited notation and choice of columns to consider. GOOG-AT-MDL-008881638 at -8. October 30, 2014. "Rethinking Bernanke: Grid search to line search."

Table 6: Table demonstrates the auction outcomes in various situations under Project Bernanke.

Outcome Under Project Bernanke	Mathematical Condition	Revenue (amount paid by advertiser to GDN)	Payout (amount paid by GDN to AdX)	Profit (Revenue – Payout)	Relative Margin (Profit/Revenue)
GDN second prices itself. Even after discounting by β, $\beta*b_2$ still beats c.	$c \leq \beta*b_2$	b_2, because the internal GDN clearing price is b_2.	$\beta*b_2$, because $\beta*b_2$ is the second-highest bid in AdX (and beats the reserve).	$(1-\beta)\ b_2$	$(1-\beta) > 14\%$. Recall that the Bernanke multiplier β is always less than 0.86, perhaps much less.
Because of Bernanke, GDN pays c to AdX instead of b_2.	$\beta*b_2 < c < 0.86$ B2	b_2, because the internal GDN clearing price is b_2.	c, because Bernanke lowers GDN's second bid below c.	$b_2 – c$	$(b_2–c)/ b_2 > 14\%$, because $c < 0.86* b_2$.
GDN pays c to AdX, with or without Bernanke	$0.86b_1 \geq c \geq 0.86\ b_2$	c/0.86, because the AdX clearing price exceeds the internal GDN clearing price.	c because GDN's second price is not high enough to raise AdX's clearing price.	$(1/0.86 – 1)*c$	14%, because everything clears as normal.
GDN pays c to AdX because of Bernanke lowering take-rate but not overbidding	$0.86b_1 \leq c \leq b_1$ – highest GDN bidder big enough to win without GDN's take-rate, but not with 14% take-rate	b_1, because GDN takes the largest margin it can while still forwarding a winning bid to AdX.	c, because GDN's second price is not high enough to raise AdX's clearing price.	$b_1 – c$	$(b_1-c)/ b_1 < 14\%$, because GDN lowered its margin to clear.

Because of Bernanke, GDN wins at price higher than the advertiser value.	$b_1 < c < \alpha b_1$	b_1, because GDN dips as little as possible into the Bernanke pool	c, because GDN's second price is not high enough to raise AdX's clearing price.	b_1-c	$(b_1\text{-}c)/ b_1 < 0\%$, because GDN needs to subsidize winner with Bernanke pool.
GDN does not win, even with Bernanke	$\alpha b_1 < c$	0	0	0	14%

Table 7: Table demonstrates the auction outcomes in various situations without Project Bernanke.

Outcome without Bernanke	Mathematical Condition	Revenue (amount paid by Advertiser to GDN)	Payout (amount paid by GDN to AdX)	Profit (Rev – Payout)	Relative Margin (Profit/Revenue)
GDN second prices itself.	$c \leq \beta^* b_2$	b_2, because the internal GDN clearing price is b_2.	0.86 b_2, because 0.86 b_2 is the second highest bid in AdX (and beats the reserve).	0.14 b_2	14%
GDN second prices itself	$\beta^* b_2 < c < 0.86\ b_2$	b_2, because the internal GDN clearing price is b_2.	0.86 b_2, because 0.86 b_2 is the second highest bid in AdX (and beats the reserve).	0.14 b_2	14%
GDN pays c to AdX, with or without Bernanke	$0.86 b_1 \geq c \geq 0.86\ b_2$	c/0.86, because the AdX clearing price exceeds the internal GDN clearing price.	c, because GDN's second price is not high enough to raise AdX's clearing price.	(1/0.86 – 1)c	14%, because everything clears as normal.
GDN does not win without Bernanke lowering take-rate	$0.86 b_1 \leq c \leq b_1$ – highest GDN bidder big enough to win without GDN's take-rate, but not	0, because GDN loses.	0	0	14%

	with 14% take-rate				
GDN does not win without Bernanke.	$b_1 < c < \alpha b_1$	0, because GDN loses.	0	0	14%
GDN does not win, even with Bernanke	$\alpha b_1 < c$	0	0	0	14%

Table 8: Table provides the comparison of the impact of Project Bernanke on publishers and advertisers in various situations.

Outcome under Bernanke	Ultimate Impact of Bernanke on Publisher (Bernanke Revenue – Standard Revenue)	Ultimate Impact of Bernanke on GDN Advertisers	Change in Bernanke Pool (Bernanke Profit – 14% of Bernanke Revenue)
GDN second prices itself. Even after discounting by β, $\beta * b_2$ still beats c.	$(\beta–0.86)b_2 < 0$, because second price goes down.	None, pays the same price.	$(0.86-\beta)b_2$, pool goes up exactly by how much publisher loses.
Because of Bernanke, GDN pays c to AdX instead of b_2.	$c – 0.86b_2 < 0$, because pay c instead of GDN second price.	None, pays the same price.	$0.86b_2 – c$, pool goes up exactly by how much publisher loses.
GDN pays c to AdX, with or without Bernanke	0, because same outcome anyway.	None, because same outcome anyway.	0, because margin is 14%.
GDN pays c to AdX because of Bernanke lowering take-rate	c, if c is equal to the reserve (*i.e.,* impression would have gone unsold without Bernanke). At least 0 and at most $c-0.86\ b_1$, if c is an AdX bidder (raises AdX payment from at least $0.86b_1$ and at most c, to c).	None, because indifferent between losing or winning and paying the bid (note: ROI goes down by winning and paying bid).	$0.86b_1 – c$, pool goes down in order to subsidize GDN with DRS. Could be more, equal, or less than how much publisher gains. It is less only if the impression would have otherwise been unsold, and equal only if GDN would have been the second price at $0.86b_1$. It is more whenever AdX's revenue without GDN exceeds $0.86b_1$.
Because of Bernanke, GDN wins at a price higher than	c, if c is equal to the reserve (*i.e.,* impression would have gone	None, because indifferent between losing or winning and paying the bid	$0.86\ b_1– c$, pool goes way down in order to subsidize GDN with Bernanke. Could be more, equal, or less than how much publisher gains. It is less only if the impression

highest GDN value.	unsold without Bernanke). At least 0 and at most c-0.86 b_1, if c is an AdX bidder (raises AdX payment from at least 0.86b_1 and at most c, to c).	(note: ROI goes down by winning and paying bid).	would have otherwise been unsold, and equal only if GDN would have been the second price at 0.86 b_1. It is more whenever AdX's revenue without GDN exceeds 0.86 b_1.
GDN does not win, even with Bernanke	0, because same outcome anyway.	None, because same outcome anyway.	0